A STUDY OF CLUTTER REDUCTION TECHNIQUES IN WIDE BANDWIDTH HF/VHF DEEP GROUND PENETRATING RADAR

by

Darien J. Hammett, Captain, USAF

A Thesis Submitted in Partial Fulfillment of the

Requirements for the Degree of

MASTER OF SCIENCE

in

Electrical Engineering

Approved by:

Professor_____
(Dr. Vincent J. Amuso - Advisor)

Professor_____
(Dr. Raghuveer M. Rao)

Professor_____
(Dr. John R. Schott)

Professor_____
(Dr. Soheil A. Dianat)

Professor_____
(Dr. Robert J. Bowman - Department Head)

DEPARTMENT OF ELECTRICAL ENGINEERING

COLLEGE OF ENGINEERING

ROCHESTER INSTITUTE OF TECHNOLOGY

ROCHESTER, NEW YORK

AUGUST 2002

Report Documentation Page

Report Date	Report Type	Dates Covered (from... to)
01AUG2002	N/A	-

Title and Subtitle	**Contract Number**
A Study of Clutter Reduction Techniques in Wide Bandwidth HF/VHF Deep Ground Penetrating Radar	**Grant Number**
	Program Element Number

Author(s)	**Project Number**
Hammett, Darien J.	**Task Number**
	Work Unit Number

Performing Organization Name(s) and Address(es)	**Performing Organization Report Number**
Rochester Institute of Technology	

Sponsoring/Monitoring Agency Name(s) and Address(es)	**Sponsor/Monitor's Acronym(s)**
The Department of the Air Force AFIT/CIA, Bldg. 125 2950 P. St. Wright-Patterson AFB, OH 45433	**Sponsor/Monitor's Report Number(s)**

Distribution/Availability Statement
Approved for public release, distribution unlimited

Supplementary Notes
The original document contains color images.

Abstract

Subject Terms

Report Classification	**Classification of this page**
unclassified	unclassified

Classification of Abstract	**Limitation of Abstract**
unclassified	UU

Number of Pages
146

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank God for His constant blessings and for granting me this wonderful opportunity. I would also like to thank my parents for always providing an encouraging environment and instilling in me the values required for success. Thank you to Dr Amuso for guiding my research and providing an atmosphere that facilitates learning and my committee for supporting me in this endeavor. Thank you to the Air Force Research Laboratory at Rome Laboratories, for allowing access to the system in which this research was based and answering my many questions. Thank you to Black River Systems for providing the necessary data used in the algorithms presented here. Last, but certainly not least, thanks to my wife for her faithful love and support.

Abstract

Reducing clutter is one of the most daunting problems a radar processing engineer faces. Clutter causes a significant problem when attempting to detect sub-surface targets, as any significant change in the ground dielectric will produce a return at the receiver. The difficulty in reducing the clutter is compounded by the fact that the spectral characteristics of the clutter are similar to that of the target.

While there are many methods that exist to reduce clutter, few do not require *a priori* information of either the target or the clutter. There are applications, of interest to the electromagnetic community, that are restricted in the amount of *a priori* information available to them. Estimation-subtraction filters calculate an estimate of the clutter from the statistics of the data collected and subtract that estimate from the original data. The Wiener filter has long been used as a way to suppress noise signals when a target reference is known. Using it to reduce clutter is a relatively new area of research. This research proposes estimation-subtraction filters and an application of the Wiener filter, which do not require *a priori* information to reduce the clutter of a bi-static synthetic aperture based, wideband deep ground penetrating radar system. The results of applying these filters to data collected in this way, at these depths, are illustrated here for the first time.

Chapter 1

Introduction

Ground Penetrating Radar (GPR) systems emanate electromagnetic energy into the ground for the purposes of imaging below the surface. They measure the reflected energy to form an image of the underlying structures. Most of the recent developments have been in detecting and imaging relatively shallow objects, no more than a few feet from the surface. Applications range from imaging underground power and water lines for the homeowner to humanitarian demining missions in third-world countries. However, there are other applications that require deep penetration systems. These include geological and archaeological surveys, as well as some military missions. In short, GPR appears to be a growing field in the radar community.

Whether designed to find anti-personnel landmines near the surface or important strategic facilities hundreds of feet deep, all GPR systems share the dilemma of dealing with the heavy clutter inherent in most ground environments. Designing appropriate algorithms to determine whether an object exists in a particular region is one of the most daunting and challenging problems a radar processing engineer experiences. The ability to suppress the contribution of a heavy clutter environment in an image scene is of particular interest to the electromagnetic community, as clutter is often the limiting factor in a radar system. This research focuses on reducing the clutter existing in a particular GPR system.

In any radar system employing electromagnetic energy, a significant change in the

dielectric properties of the medium of propagation causes a reflection and a refracted

transmission of the illuminating wave front. Whenever electromagnetic energy impinges

on a boundary, some energy is transmitted through the medium and some energy is

reflected away. The transmitted energy follows Snell's Law and the reflected energy

abides by the Law of Reflection. Snell's Law states that when electromagnetic energy

encounters a medium that differs in refraction index from the current medium, the

propagation direction is changed by an angle that is defined by the ratio between the

incident medium's refraction index and the refractive medium's refraction index. The

Law of Reflection states that the angle of incidence will equal the angle of reflection,

where the angles are defined between the direction of propagation and the surface

normal. Figure 2.1 is an illustration of Snell's Law and the Law of Reflection where the

incident energy is blue, the reflected energy is red, and the transmitted energy is green.
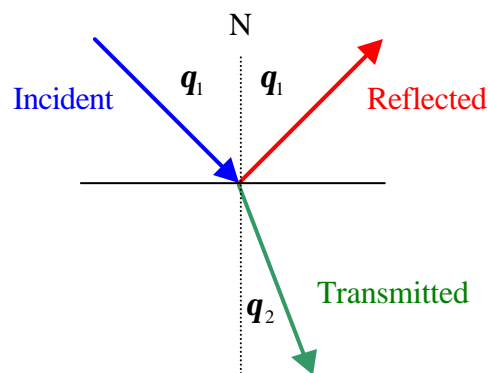
Figure 1.1 Reflected and Transmitted Energy

Clutter is described as any unwanted radar return [1, 2]. Since both a target and

sources of clutter have the potential of reflecting the incident wave, the task of this work

is one of determining whether the echo is from a target or clutter. GPR clutter can take

many forms, some of which are tangible and some not. Tangible sources of clutter

include undulations in the earth's surface, vegetation, water, and rock. What they all

share in common is that they represent a change in the dielectric permittivity of the

propagation medium. An intangible source of clutter would be the direct path signal

from the transmitter to the receiver.

It is certainly desirable to eliminate the clutter as much as possible, because the

unwanted returns obscure the target and can obviously degrade a systems ability to detect

the target. Also of the causes that limit the performance of ground penetrating radar,

clutter is usually the dominating factor [1, 3]. In this particular system, it is assumed that

clutter is the limiting factor for detecting the target. In Schott's application of the image

chain approach, he states that making improvements to the limiting factor can provide the

most benefit to the entire system [4]. Consequently, clutter limited processing will be

used and noise limitations will not be examined.

1.2 Clutter Reduction in Radar

There are many signal processing methods used to reduce clutter. The difficulty

in reducing clutter from a radar scene stems from the fact that the spectral density of

clutter is very similar to that of the target [1, 2, 5, 6]; therefore, the clutter return

resembles the target return. It is in this way that the clutter is different than noise. Noise

and clutter both degrade the target signal (and are therefore undesirable), but while noise

is random in nature and generally differs spectrally from the target return, the clutter

spectral density is similar to that of the target. Conventional radar systems can mitigate

clutter by using moving-target-indication techniques (MTI). However this problem is

exacerbated in GPR systems, because the scenes are static.

Whenever detection of a particular object is the specific goal of a GPR system,

the users would like to collect as much *a priori* information about the target and

environment as possible. This information would enable the selection of filter parameters

or situation specific algorithms to combat that distinct problem. Such things as the

reflectivity functions of the target object and environment, the composition of the ground

itself, and size and shape of the target would be valuable pieces of information to assist

the system in detecting the specified target. However, there are many applications where

this information is extremely difficult to get, or simply not available. Not having *a priori*

information complicates the clutter reduction problem.

This research employs classical image processing techniques to reduce the clutter

of three-dimensional GPR scenes. Some adaptations were made to the classical

techniques so that their application was appropriate for the specific situation of interest.

Some of the adaptations stem from a self-imposed constraint of *a priori* information

being unavailable. The reasons for this constraint revolve around the intended uses of the

system providing the data. An overview of the remainder of this thesis follows.

Chapter 2 reviews some of the previous work in the area of clutter reduction applied to
GPR systems.

Chapter 3 walks through the system hardware and processing to which the clutter
reduction algorithms were applied.

Chapter 4 introduces the theory of the estimation subtraction and Wiener filter, which the algorithms are based.

Chapter 5 discusses the details of, as well as the adaptations made to, the image processing algorithms employed.

Chapter 6 presents and discusses the results of the algorithms applied to real three-dimensional data sets.

Chapter 7 submits conclusions and suggests areas of future work.

Chapter 2

Literature Review

Researching previous work is necessary in order to take advantage of the prior advances of others. Viable solutions must be analyzed before selecting the optimal one for the specific conditions that exist in a specific application. A discussion of five general methods of clutter reduction follows, along with justification for selecting the options pursued. In chapter 3, the GPR system to which these clutter reduction algorithms are being applied will be examined. The background theory for these options is discussed in chapter 4. A discussion of the details and changes made to the techniques associated with these options occurs in chapter 5.

2.1 System Identification

The first step in representing an imaging system by System Identification is to break the image chain up into blocks, representing every element in the system, before modeling each block with a transfer function. Although Schott's application of system identification is in remote sensing [4], he uses the transfer function to model each link from the image chain as a filter that alters some range of frequencies of interest. He then exploits this information to determine where to focus the attention necessary to improve image fidelity. In Schott's image chain representation, clutter can be treated as just one of links in the image chain that contributes to the degradation of an image. Brooks discusses system identification applied to GPR signals [2]. A simplified system identification model is shown in Figure 2.1 [2]. From this diagram an equation for s_{out}

can be found as,

$$s_{out} = s_{in} * h_{a1} * h_{c1} * h_t * h_{c2} * h_{a2} \qquad (2.1)$$

where s_{out}, s_{in}, h_{a1}, h_{c1}, h_t, h_{c2}, and h_{a2} are the output signal, input signal, and the impulse

responses of the transmit antenna, clutter between the transmitter and the scatterer, target,

clutter between the scatterer and receiver, and receive antenna, respectively. The symbol
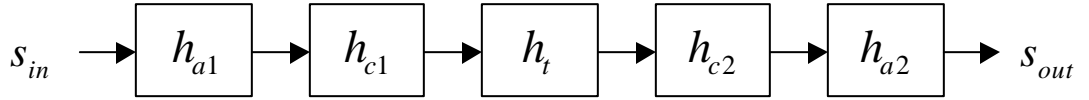
Figure 2.1 Simplified System Identification Model

"*" represents linear convolution, which is defined as

$$x(n) * h(n) = \sum_{k=-\infty}^{\infty} x(n-k)h(k) \qquad (2.2)$$

for discrete sampled data.

Modeling of a system in this manner requires that the system be linear. System

Identification attempts to estimate the transfer function of every element of the image

chain. In many links of the image chain, such as the physical antenna or optics, this can

be accomplished in the laboratory by finding the impulse response [4]. Under certain

circumstances, estimating the target and the clutter impulse response can also be

accomplished in the lab. However, it is usually more difficult, because of the variance of

the target and clutter impulse responses with respect to aspect angle and soil conditions

[1, 2]. In many cases, the transfer function of the clutter can be found by taking a target-

free sample in the field environment, or by taking samples of the clutter and duplicating

the field environment in the lab. Both of these methods were performed successfully by

Sahli et al in Angola while performing humanitarian demining missions.[1]

2.2 Recursive Filtering

Recursive methods are adaptive in nature and usually applied in real-time

processing of data as opposed to off-line processing [1, 2]. Real-time processing has the

mandatory condition that processing of measurements from one sample must be

completed during one sample interval [7]. In recursive filtering, clutter is estimated

based on past values of the clutter. The estimation of clutter can take many forms.

Brooks et al discuss a recursive least square method where a target-free sample of clutter

is required and a linear predictive coding (LPC) method that does not require any *a priori*

information. Both were used in a landmine-detecting GPR system to reduce clutter

caused by changing undulations of the surface [1]. A target-free sample of clutter was

found by sampling the environment that was known not to contain any targets. Brooks

used LPC to estimate the prediction-error coefficients of his linear regression model,

applying the Levinson-Durbin algorithm to the autocorrelation matrix of the input data.

Once the input data is filtered with the prediction-error coefficients, the filtered estimate

is subtracted from the input data to produce the final output data of interest.

2.3 Matched Filtering

The detection of a known input signal embedded in noise using a matched filter is

shown by Skolnik [8]. The matched filter performs this function well because its

response maximizes the output peak-signal-to-mean-noise power [8]. A basic matched

---

[1] Per conversations between the author and Dr Hichem Sahli of Vrije Universiteit Brussel.

filter design is given in Figure 2.2. Assume a point target located at some distance, *d*,

from the source is illuminated by the transmit signal, *s(t)*. Upon encountering the target,

the transmitted signal, *s(t)*, plus noise, *n(t)*, returns and is denoted as the echo signal, *r(t)*.

The echo is transformed to the frequency domain via the Fourier Transform, *R(f)*. This

signal is then mixed with the Fourier Transform of a delayed conjugated version of the

original reference signal, *S*(f)*. The time delay, *T*, is a fixed value of time at which the

signal is observed to be a maximum. The product resulting from multiplying is then

inverse Fourier Transformed to form *s^(t)*. The target is detected when *s^(t)* exceeds a

pre-determined threshold.
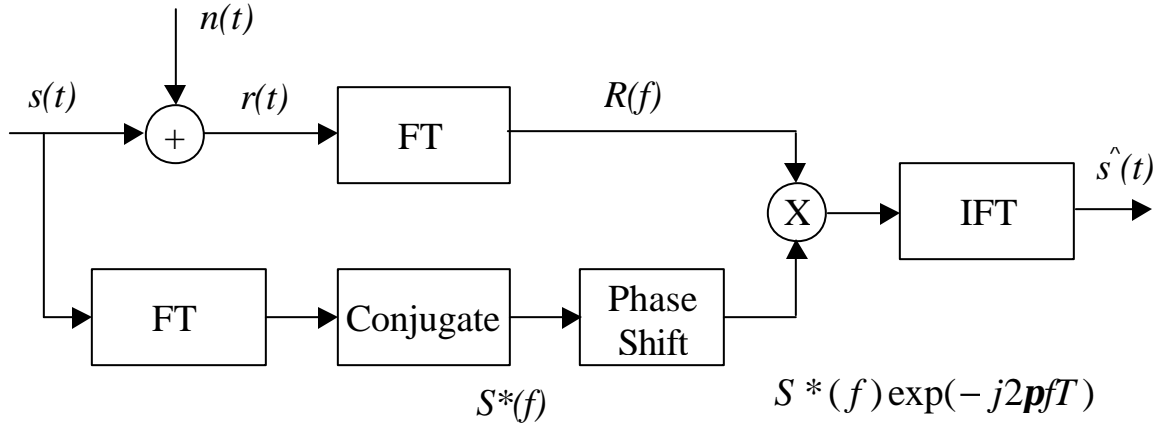
Figure 2.2 Matched Filter Basic Design

The continuous Fourier Transform is a method to transform a time domain signal,

f(t), into the frequency domain. Although the operations of a matched filter could be

performed in the time domain, that would require the use of a convolution operation.

Performing this process in the frequency domain allows us to use multiplication, which is

9

much simpler.

The minimum noise and correlation energy (MINACE) filter is a matched filter that was originally designed as an optical filter to produce easily detectable correlation peaks in the presence of noise and clutter [9]. Ravichandran and Casasent [9] accomplished this by combining the ideas of the minimum variance synthetic discriminant function correlation filter (which minimizes noise variance in the output correlation peak/plane) and the minimum average correlation energy (MACE) filter (which minimizes the average correlation plane energy over all training sets). The MINACE filter transfer function is given by [9]

$$\mathbf{H} = \mathbf{T}^{-1}\mathbf{X}(\mathbf{X}^{/}\mathbf{T}^{-1}\mathbf{X})^{-1}\mathbf{u} \qquad (2.3)$$

where $\mathbf{T}$ represents a positive definite weighting matrix, $\mathbf{X}$ is a matrix composed of a set of target returns from different aspect angles, and $\mathbf{u}$ is a vector of the correlation constraint peak values. The symbol "/" represents the conjugate transpose operation. The MINACE filter works well provided *a priori* target information is available. However, in practice they require a large database of filters containing *a priori* information from many aspect angles in order to be productive [9, 10, 11]. In 3D space the number of aspect angles required can grow impractically large, because of the many different aspect angles that must be accounted for in a complex target. Requiring this much *a priori* information may not always be practical. The performance of the filter is dependent on the amount of *a priori* information regarding different aspect angles.

2.4 Wiener Filtering

In signal processing, the Wiener filter is well known as a means to discriminate a

reference from noise. It is an optimal filter in the sense that it minimizes the mean

squared error between the restored image and the original [12, 13, 14]. This assumes that

the reference and noise signals are uncorrelated, stationary random processes and that the

noise is a zero-mean process. Its use has recently been expanded in optical pattern

recognition systems to include detection of targets in unknown clutter environments.

Marom and Inbar [12] combine a restoration Wiener filter with an inverse filter and apply

it to two-dimensional images. They illustrate the effectiveness of the filter by identifying

a simulated reference object out of various noise environments. Tan et al [13] apply the

same filter to optical images with limited *a priori* knowledge of the clutter environment,

but have a known reference. In order to make the simulation more realistic, they defined

the clutter reference with the average of multiple backgrounds. The main observation

from their results is that the peak to correlation energy values drop rapidly when the

clutter backgrounds are not the same as those used to construct the Wiener filter. One

conclusion noted by the authors was that the Wiener filter was very sensitive to the

clutter term and that detection of the target was dependent on how well the scene

background was matched by the filter clutter term. Therefore, an image with a

background different from what was used to create the filter reference does not produce

reliable detection of the target in the scene. A possible solution to this problem is to

estimate the clutter power spectral density with the power spectrum of the entire scene.

Although Tan et al suggested this technique, neither they nor Marom and Inbar applied it

to their optical images, let alone radar images.

Estimation-subtraction is a process whereby the clutter is estimated in some fashion from the input signal and subsequently subtracted out of the input signal. Lundberg [5] discusses an infrared (IR) ground penetrating landmine detection system that was subject to severe surface clutter. His method of reducing the clutter culminated in adding a sensor that operated in the visual wavelengths to the same platform that contained the existing IR sensor. The clutter in the IR image is then considered to be a linear filtered version of the data collected from the optical sensor. The clutter is then reduced by estimating the filter using an adaptive least-squares algorithm on the visual wavelength data and then subtracting the estimated clutter from the original IR image. Lundberg noted that the optical sensor was successful at capturing the uneven enlightenment of the surface, which corresponds to the warm and cool regions on the surface (the clutter the IR sensor experiences). This enabled him to remove the influence of the surface structure, improving target detection by decreasing the amount of falsely detected targets.

Despite the difference in wavelength, the idea of estimation-subtraction was the inspiration for the clutter reduction schemes applied to radar signals in this thesis. This method was chosen because it does not require *a priori* knowledge of either the target or the clutter. In addition, an application of the Wiener filter is implemented. Under most commercial or favorable situations, choosing a method that uses *a priori* information about the target or clutter would be preferred. Intuitively, the more data acquired about the exact scenario and environment surrounding that scenario, the more adequate the

means to delineate between a target and the background. However some scenarios make this difficult, if not impossible. For instance, in military applications it would be nearly impossible to predict every conceivable adversarial target clutter environment or be granted access into these denied areas to obtain such information. As will be mentioned in chapter 3, the ultimate goal is to use this GPR system in applications where *a priori* information of either the target or clutter isn't necessarily available. To this end, implementing algorithms that did not require *a priori* information about either the target or clutter environment was desired. By relying only on the data collected, application of these estimation-subtraction and Wiener filters may be utilized across a wider spectrum of targets and clutter. This approach, as well as the specific techniques used in the filters, will be discussed in detail in the following chapters.

Chapter 3

System Description

To a great extent, system parameters as well as hardware configurations and processing methods, can determine any limitations on post-processing techniques used for enhancement in an imaging system. A quick perusal of the many types of radar and processing schemes [8, 15] will convince the reader that the specific mission of a radar system has a great impact on the selected processing. In this chapter, the GPR system on which this thesis project is based will be defined. The first goal is to describe the data collection and signal processing that occurs immediately before clutter reduction is attempted. In this way, the genesis of the data and the path it takes before clutter reduction can be fully understood. This understanding will allow the reader to fully appreciate the methods that were chosen to reduce the clutter.

3.1 The GPR System of Rome Laboratories

Rome Laboratories (of Rome, New York), under the Sensors Directorate of the Air Force Research Laboratory of Wright-Patterson Air Force Base, Ohio, developed a GPR system using a Synthetic Aperture Radar (SAR) approach [16]. Currently the development is in a relatively early stage, where the antennas are placed directly on the ground pointed into the earth. The ultimate published goal of Rome Laboratories is to eventually establish the system in some type of airborne platform. It should be noted that the primary targets are buried structures of strategic importance [16]. The current system can be broken up into three stages: the hardware or data collection system, pre-

processing, and SAR processing. These stages will now be described in detail.

3.2 Data Collection System

The system hardware contains a single transmitter and multiple receiver locations that form the synthetic aperture. A mixer circuit, power supply, analog-to-digital converter, arbitrary waveform synthesizer (AWS), and personal computer (PC) comprise the hardware platform, with coaxial cables throughout the system. A block diagram of the system hardware is given in Figure 3.1.
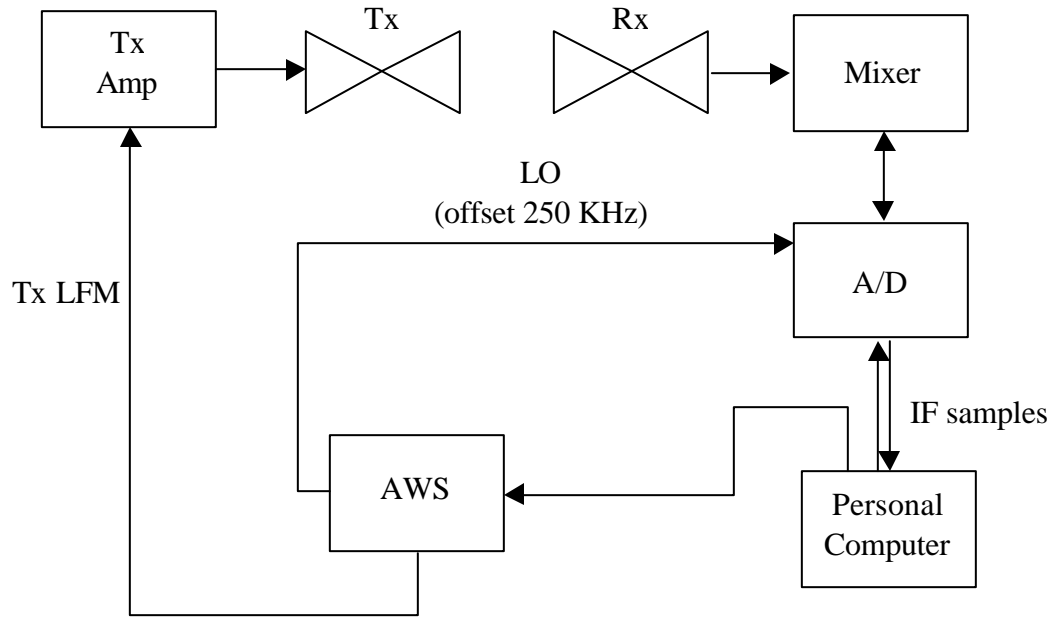
Figure 3.1 Data Collection Block Diagram

3.2.1 Personal Computer

A custom built personal computer manages the entire data collection operation. The PC was assembled with a Supermicro P6SLA motherboard and a Pentium II 300

MHz processor. It sends a trigger pulse to the Arbitrary Waveform Synthesizer (AWS) to generate the waveform, which begins the collection process. A collection period is defined as the time that the transmitter is transmitting and a particular receiver is receiving and lasts approximately two seconds. Between 900 and 1200 pulses are transmitted during the collection period of approximately two seconds. After the collection period, each receiver's pulse data is stored as a file on the 10 GB hard drive of the PC.
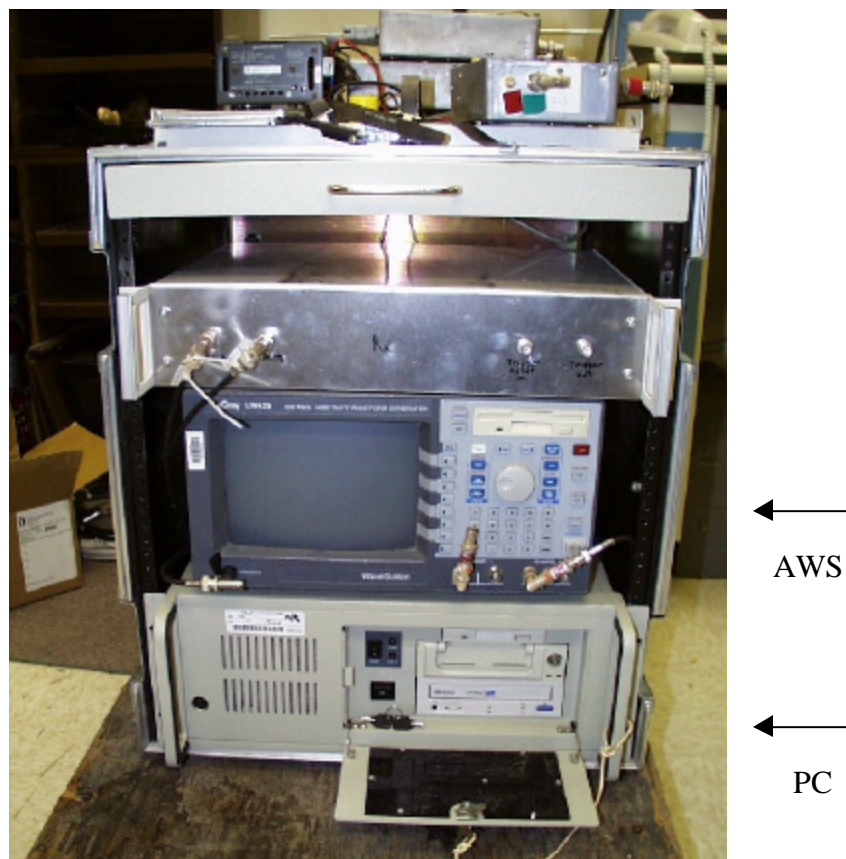
Figure 3.2 AWS (top) and Personal Computer (bottom)

3.2.2 *Arbitrary Waveform Synthesizer*

The AWS creates the signal that is transmitted and the local oscillator (LO) signal

that is sent to the mixer, which is used for down-converting the received pulse. The

AWS is a Lecroy LW420 8-bit digital-to-analog converter (Figure 3.2), which samples

the data at 400 MHz. The pulse created for transmission is a two-millisecond down-

chirped linear frequency modulated (LFM) signal with a bandwidth from 66 MHz to 6

MHz. Figure 3.3 gives an example of an LFM chirp waveform. The pulse is linear

weighted giving more amplitude to higher frequencies than lower frequencies. This

weighting is performed to account for losses the pulse experiences as it is propagated

through the coaxial cables. The local oscillator (LO) reference signal is the same as the

transmitted pulse with a 250 KHz offset. This offset was selected for two reasons. The

first reason is that the 250 KHz offset provides an intermediate frequency (IF) that is far

removed from low frequency noise such as that produced by switching power supplies.

Also, A/D sampling at 1 MHz simplifies the digital conversion from an IF of 250 KHz to

baseband (section 3.3.2).
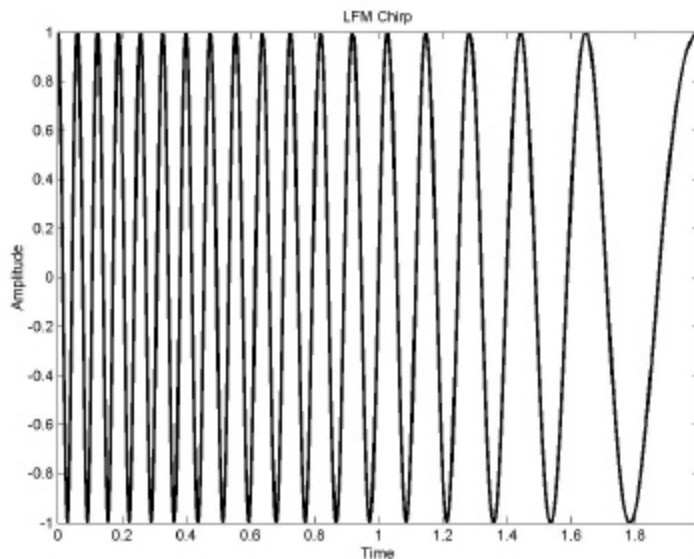
Figure 3.3 Example of an LFM Chirp

3.2.3 Transmitter

The transmitter consists of a stationary bow-tie antenna and an amplifier. The bow-tie antenna is 22-feet by 5-feet and is constructed with galvanized conduit (Figure 3.4). The bow-tie topology was chosen because of its large beamwidth. Although a narrow beamwidth would normally be desired for resolution purposes, a large beamwidth is actually desired in this case because then only one transmit antenna is required to illuminate the entire suspected target area. This is acceptable, because the resolution will be improved by using multiple receiver locations in forming the synthetic aperture. This technique will be discussed in detail later in section 3.4. The amplifier amplifies the outgoing signal by 50 dB. The transmit antenna transmits the signal generated by the AWS after amplification.

Figure 3.4 Transmitting Antenna (right) and Receiving Antenna (left)

3.2.4 Receiver Array

To create the synthetic aperture, the receiver measurements are taken along a grid of receiver locations separated by 10-feet in each dimension. The most common array used in this system consisted of a square grid with a size of 100-feet by 100-feet (although smaller arrays and arrays as large as 100-feet by 300-feet have been used). This results in 121 receiver locations on the ground at a spacing of 10-feet in both dimensions (Figure 3.5). The receiver array is generated by physically moving a single
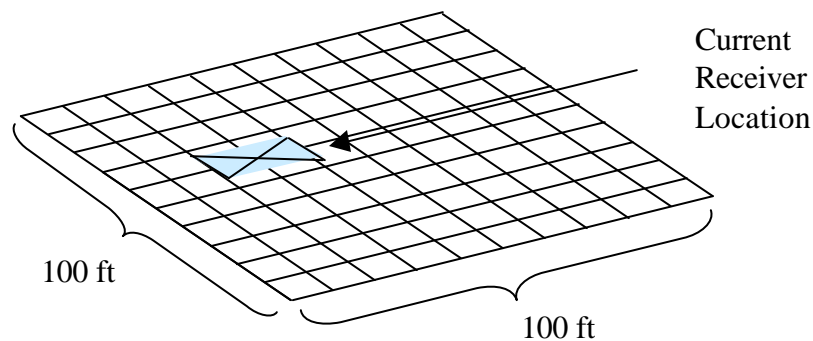
Figure 3.5 Receiver Array

receiving antenna to all of the receiving locations. The antenna is placed in position for the data collection period, before the PC triggers the beginning of the collection period and transmission commences. Once the data collection period has ended, the receiver is moved to the next location and the process is repeated. This continues until the receiver has been positioned and collected data at all receiver locations.

The physical receiver antenna is very similar to the transmit antenna. It is a 14-feet by 3-feet bow-tie antenna that is also made from galvanized conduit. It is slightly smaller than the transmit antenna for ease of relocating it after each collection period

(Figure 3.4).

3.2.5 Mixer

The mixer circuit consists of a radio frequency (RF) amplifier, a mixer, an LO amplifier, a low pass filter (LPF), and an IF amplifier and is attached to the back of the receiving antenna (Figure 3.6). The RF amplifier amplifies the incoming signal by 15
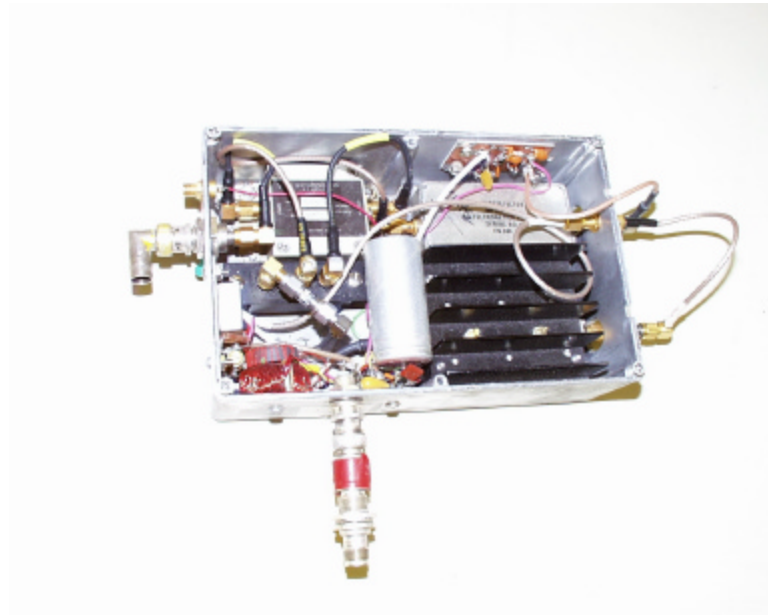
Figure 3.6 Mixer Box

dB. The mixer is a Lorch FC-245Y high-level mixer which multiplies the incoming amplified received signal with the LO reference signal to remove the modulation of the received signal. After passing through the mixer, there will be a sum and a difference component of the original signal. To complete the down-conversion process the LPF selects the difference component. Since the cut-off frequency of the LPF is set to 1 MHz, in addition to selecting the difference component, it prevents any RF (6-66 MHz) leakage from corrupting the IF (~250 KHz) video signal. The signal now exists as a

superposition of all the individual demodulated reflections received at this location. The time-delay of the pulse corresponds to the tone produced by the mixer (Figure 3.7). When mixed, the reflected pulse (shown in black in Figure 3.7) and LO (shown in red in Figure 3.7) produce a tone directly related to the location of the scatterer that created the reflection. In the low frequency example of Figure 3.7, the LO is mixed with an echo that has a small time-delay with respect to the transmitted signal and the frequency difference is small. In the high frequency example of Figure 3.7, the LO is mixed with an echo that has an increased time-delay and the frequency difference is large. The time-delay of the echo is directly related to the distance traveled by the pulse from the transmitter to the scatterer that produced the reflection detected by the receiver. A particular receiver will detect echoes from a number of scatterers in the scene, so that
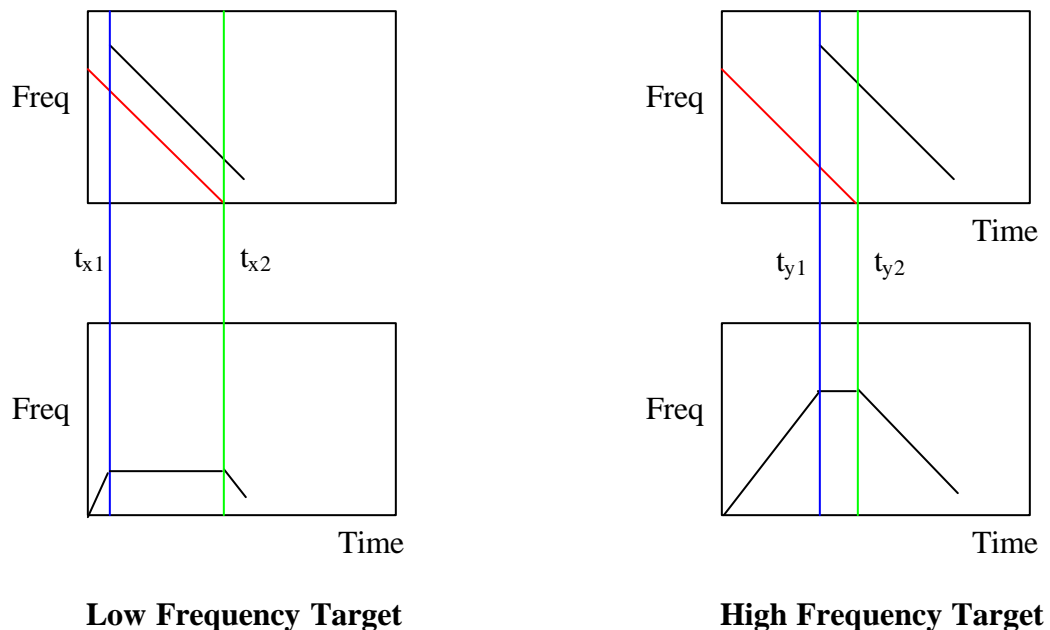
Figure 3.7 Time-Delay Relationship to Target Tone

each mixed signal is now comprised of a superposition of tones. The IF signal is then amplified by a factor of 30 dB by an H0550 Analog Devices IF amplifier, which is routed to the A/D electronics.

3.2.6 Analog-to-Digital Converter

A Datel ADS-937 A/D converter samples the IF analog signal from the mixer circuit so that it can be stored in the PC and processed off-line. It is shown in Figure 3.8. The sample frequency was chosen to be 1 MHz, which is four times the sampling rate used in the I-Q Demodulation. It will be shown in section 3.3.2 that this simplifies the digital conversion to baseband. For the pulse length of 2 ms used in this system, the sample frequency results in 2000 samples/pulse. After sampling, the data is latched into a FIFO as 16-bit parallel data and stored. At this point the data is referred to as "raw" and will ultimately be handled by the pre-processor. The raw data consists of the 900-1200 pulses collected per receiver location; each individual receiver data set is stored as an individual file.
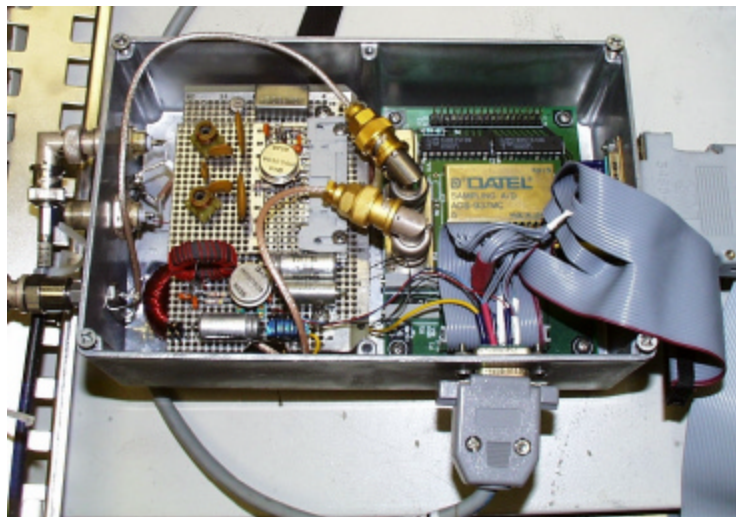
Figure 3.8 A/D Converter

22

Power is supplied to the hardware components by means of a Lambda LDS-W-28

power supply (Figure 3.9). The power supply has a maximum output of 28 volts at 4.7

amps and integral power regulators provide the specific digital and analog voltage levels

needed for particular devices. The PC has a separate power supply. A 4.5 kW Honda

gasoline generator provides the power for all of the equipment hardware.
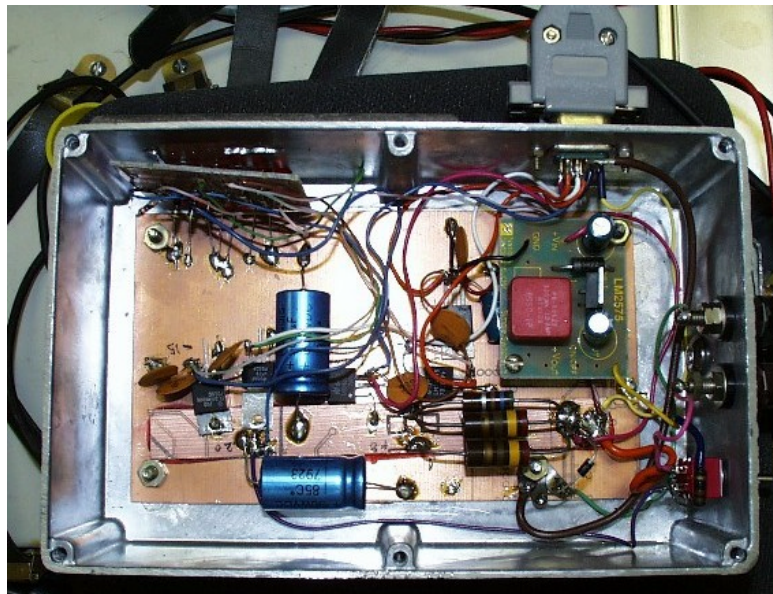
Figure 3.9 Power Supply

3.3 Pre-Processing

Before the data from the separate receiver locations can be aligned to register the

receiver returns to specific voxel (volume-pixel) locations in the scene and formed into

an image, the data must be pre-processed. The pre-processing block diagram is

illustrated in Figure 3.10. The ensemble of pulses for a given receiver location are
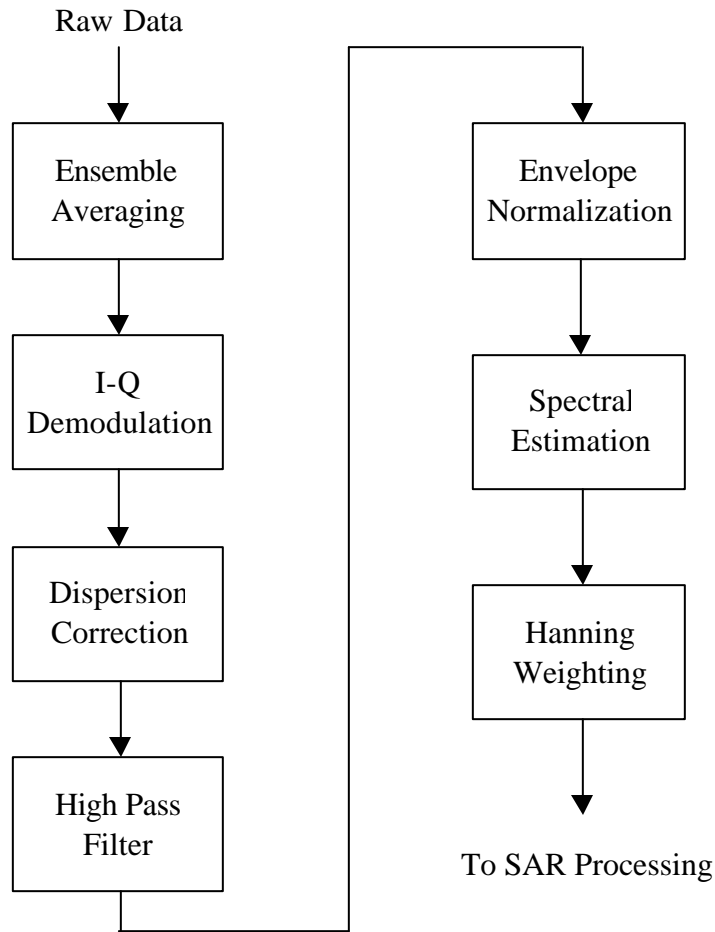
Figure 3.10 Pre-Processing Block Diagram

averaged and then split into the in-phase (I-channel) and quadrature-phase (Q-channel)

components. Following I-Q demodulation, the pulses are corrected for frequency

dispersion and high pass filtered. The envelope of the resulting data pulse is then

normalized and spectral estimation is performed. Hanning weighting is finally applied

leading to the data that is forwarded to SAR Processing. The detail of these steps

follows.

3.3.1 *Ensemble Averaging*

The first step of pre-processing is to average every ensemble of the 900-1200 pulses received by each individual receiver into one representative pulse by calculating the mean of all of the samples in a specific time instance for a given receiver (Figure 3.11). In this way, each ensemble for a specific receiver location is reduced to one pulse of data.

$$1x10^4 \begin{bmatrix} 3.3471 & 3.2045 & 3.1956 & \cdots \\ 3.3241 & 3.2982 & 3.2161 & \cdots \\ 3.2295 & 3.3103 & 3.3113 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \longrightarrow 1x10^4 \begin{bmatrix} 3.3002 & 3.2710 & 3.2410 & \cdots \end{bmatrix}$$

Figure 3.11 Ensemble Averaging

An example of an averaged pulse from a single receiver location is shown in Figure 3.12. It was collected at the U. S. Department of Energy's Nevada Test Site, located northwest of Las Vegas. This pulse will be displayed following the explanation of every pre-processing step in order to demonstrate the effect these steps have on the data.
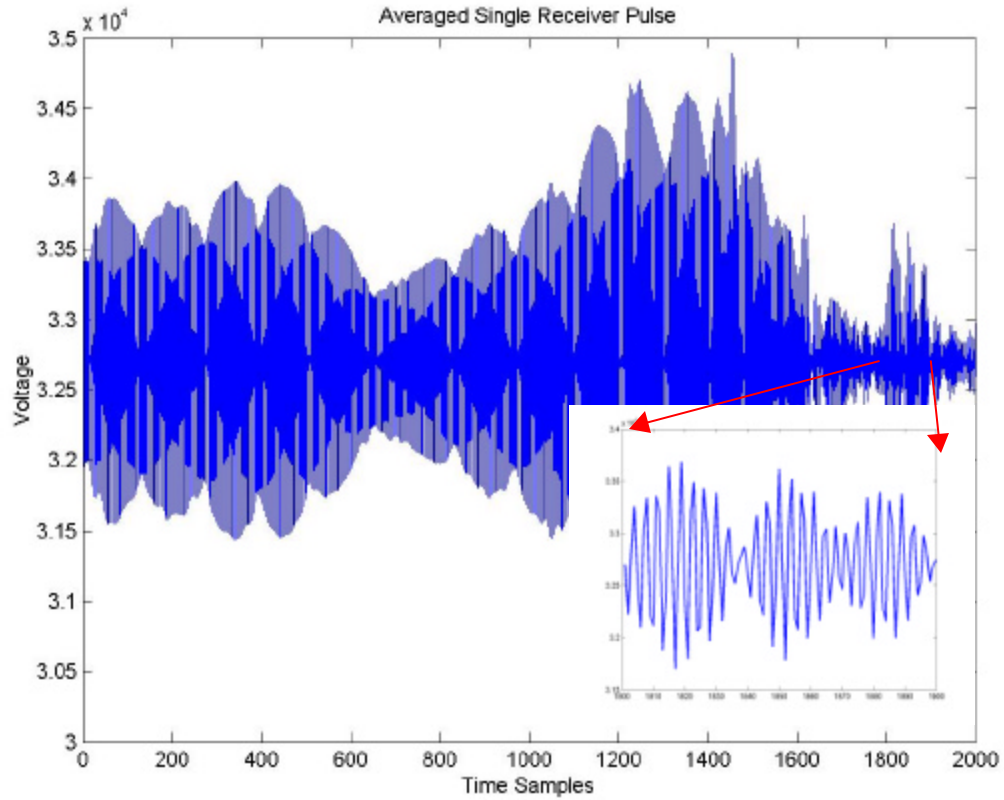
Figure 3.12 Ensemble Average for one Receiver Location

3.3.2 I-Q Demodulation

I-Q demodulation involves processing the averaged raw data from each receiver location in a manner that will facilitate the instantaneous measurement of the phase of the original signal. The I-Q demodulation procedure down-converts the over-sampled IF data to baseband. The I-channel represents the real component of the complex signal and is found by mixing the input with the cosine of the LO frequency. The Q-channel represents the imaginary component of the complex signal and is found by mixing the input with the sine of the LO frequency. Therefore, the I and Q channels are 90 degrees

out of phase from one another. Once the individual I and Q channels are produced, the

phase at a particular point in the data can be determined by calculating the arctangent of

the imaginary component divided by the real component.

In this system, the demodulation is accomplished by multiplying the I-channel

component by a sequence of +1, 0, -1, 0, +1… The Q-channel component is

demodulated by multiplying by a sequence of 0, +j, 0, -j, 0, and so on. The I and Q

channels can be processed in this straight-forward manner, because of the over-sampling

rate of the A/D converter discussed in section 3.2.6. The I-channel component is the

current sample and the Q-channel component is the following sample. The sampling rate

is 250 KHz to support a signal bandwidth of 125 KHz. This effectively decimates the IF

data by a factor of four, which results in 500 samples of data from the original 2000. In

order to prevent any high frequency aliasing associated with the decimation, a low pass

filter is applied as part of the procedure. Figure 3.13 shows the I-Q demodulated output

of the input pulse from Figure 3.12.
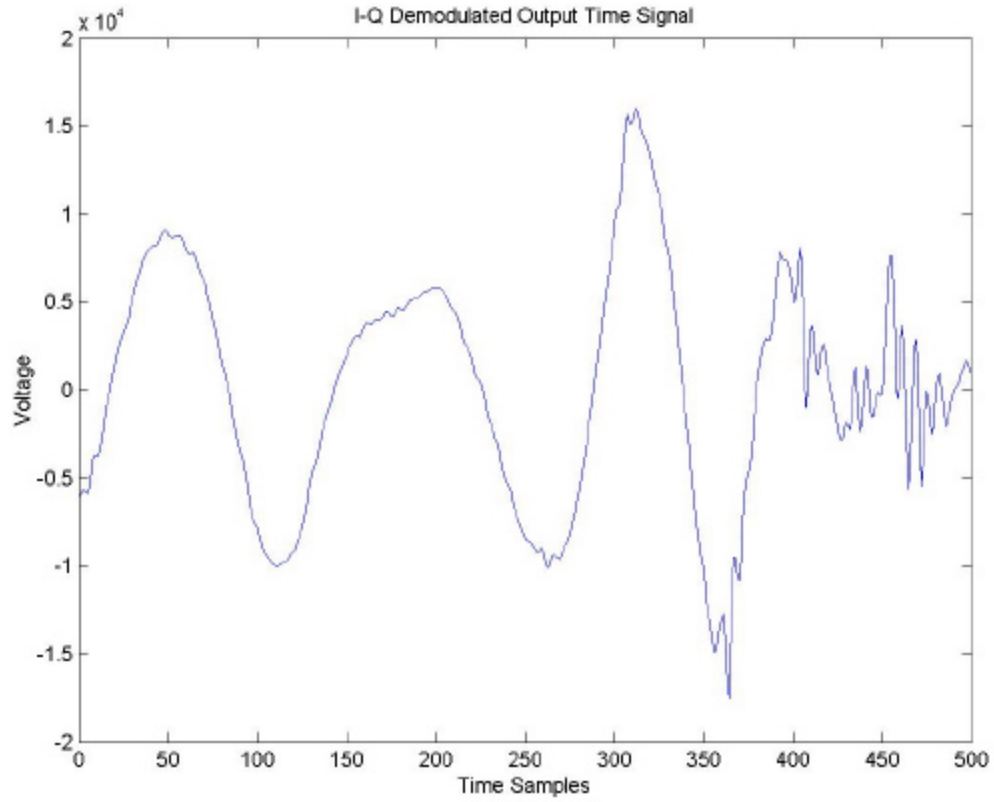
Figure 3.13 I-Q Demodulated Output

3.3.3 Dispersion Correction

As a signal propagates through a medium, different frequency components of that signal travel at different velocities through the medium. This phenomenon is referred to as dispersion. In the visible region of the electromagnetic spectrum, dispersion manifests itself by separating the different colors of white light as the light is passed through a prism [17]. In this GPR system dispersion manifests itself by producing a chirp following the mixing operation, instead of a tone. Multiplying by a phase correction factor compensates for dispersion in this system.

The phase correction was implemented using a heuristic approach. The direct path returns from a line of receivers aligned with the transmitter were analyzed following the mixing operation. Instead of observing the expected direct path tones, it was noted that another chirp was present. A phase adjustment was applied until the approximate expected frequency tone was generated at each receiver location. This phase adjustment was then adopted as the dispersion correction algorithm and has been verified with data collected following its implementation. This phase correction corrects for any dispersion experienced in both transmission and reception. Figure 3.14 illustrates the peak sharpness that occurs in the spectrum of the signal following dispersion correction.
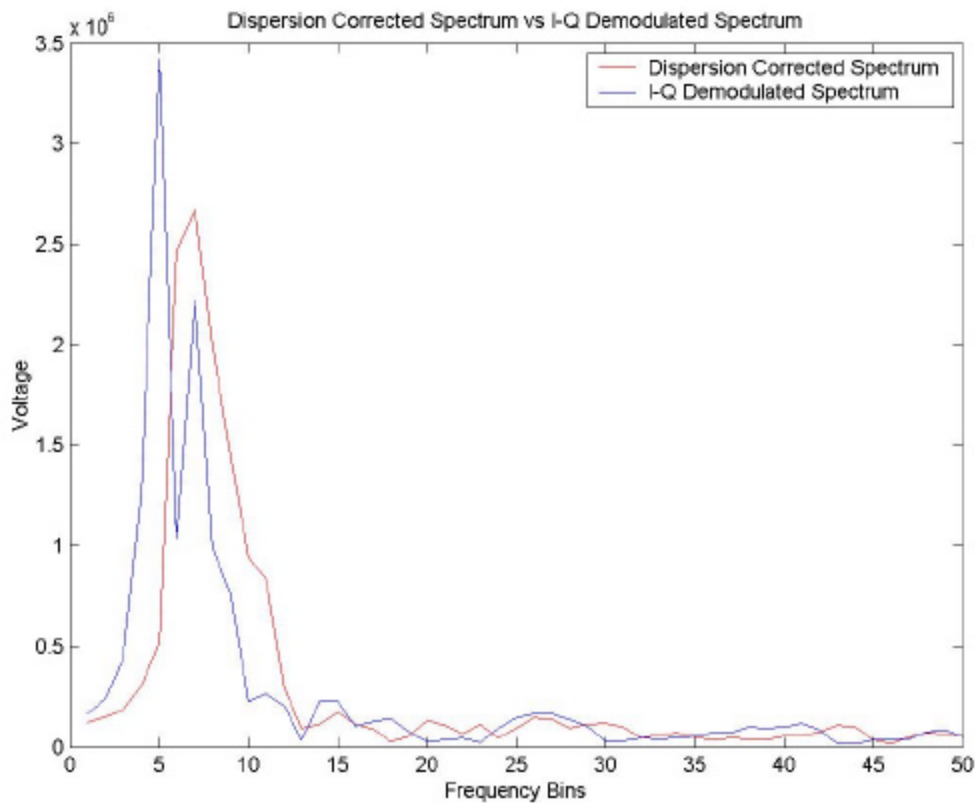
Figure 3.14 Dispersion Corrected Spectrum vs IQ Demodulated Spectrum

The high pass filter (HPF) is designed to remove the direct path signal. The direct path signal is a pulse that is sent directly from the transmitter to the receiver without reflecting off of a scatterer (Figure 3.15). The distance traveled by this signal will be shorter than a signal that reflects off of a scatterer. Therefore, the phase resulting from this direct path signal will be small and when it is mixed with the reference the resulting signal will have a low frequency. Referring to Figure 3.7 and 3.15, the pulse that traverses the path depicted by "B" will result in a high frequency tone and the direct path pulse that traverses the path depicted by "A" will result in a low frequency tone. The cut-
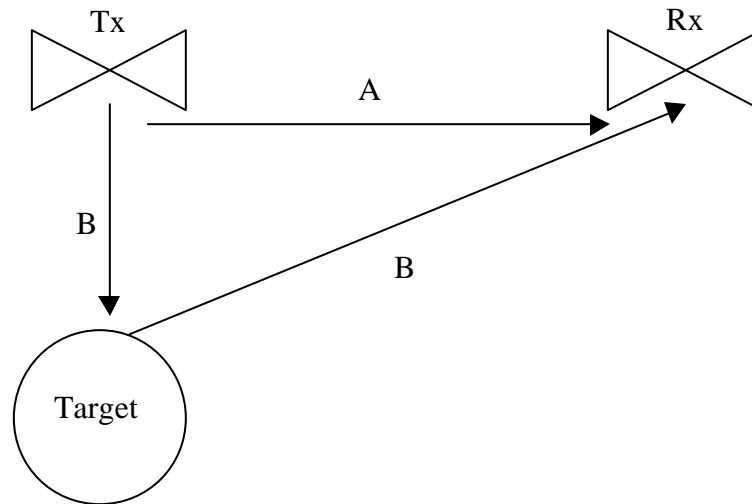
Figure 3.15 Direct Path (A) and Reflection from a Target (B)

off frequency is dependent on the direct path length from the transmitter to current receiver location. The filter coefficients are changed appropriately as the receiver location changes, because the size of the array has a significant effect on the distance of

the direct path. The HPF has its parameters set such that pulses within +/- 5 feet from the

expected direct path are eliminated. This will attenuate any signal received within a 5-

feet radius of the expected direct path for a specific receiver location. Figure 3.16

compares the time-based result of the dispersion corrected output to the HPF output.
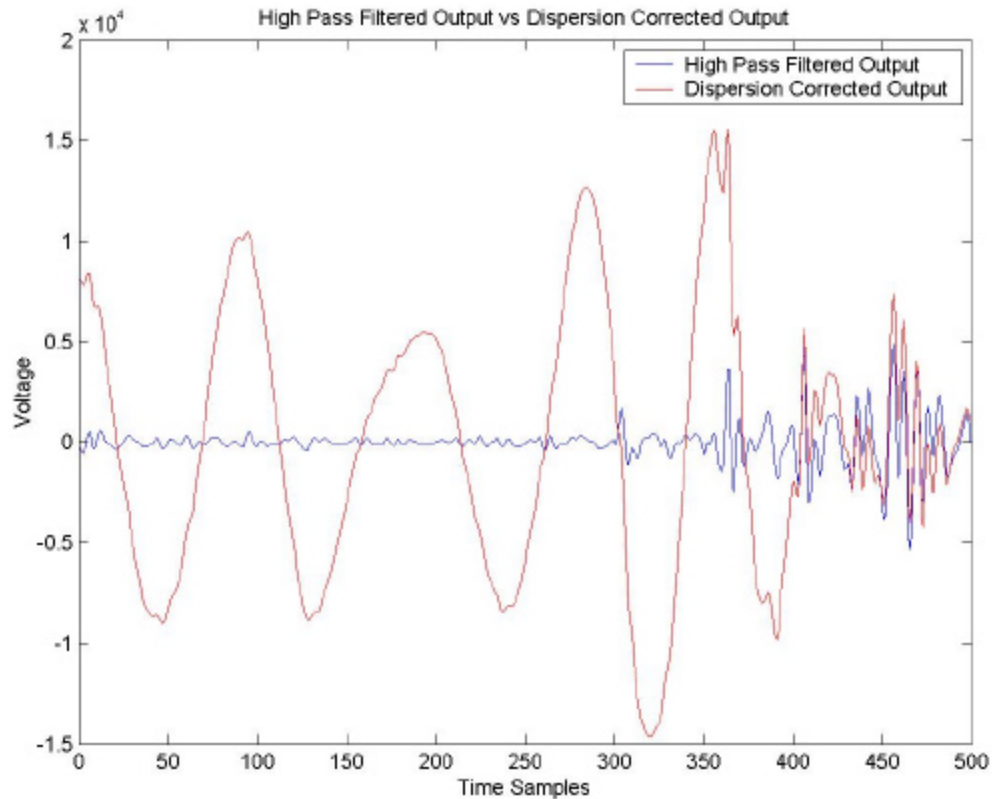
Figure 3.16 High Pass Filtered Output versus Dispersion Corrected Output

3.3.5 Envelope Normalization

As a pulse travels through a medium of propagation, different frequencies

experience different amounts of attenuation. When the ground is the medium of

propagation (as in GPR systems), higher frequencies are generally attenuated more than

lower frequencies. Envelope normalization reshapes the envelope of the incoming pulse

so that it is does not display the effects of the unbalanced ground attenuation. This is

accomplished by first calculating a fifth-order fit to the envelope of each pulse that has

been high pass filtered. The mean of all of the envelopes is then calculated. Multiplying

each high pass filtered pulse by the inverse of the mean envelope normalizes the

incoming envelope. The normalization function is shown in Figure 3.17. Thus, the

normalization function used on the incoming signal is a product of the specific
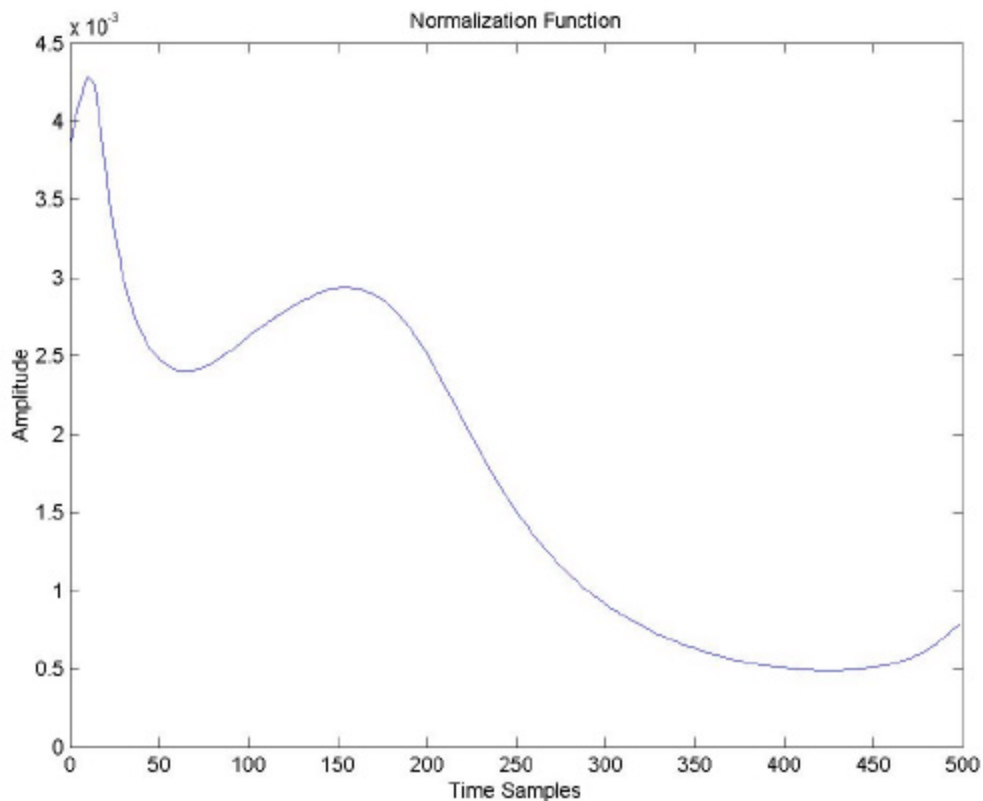
Figure 3.17 Normalization Function

attenuation effects of the propagating medium. Figure 3.18 illustrates the effect of

normalization on the HPF data shown in blue of Figure 3.16.
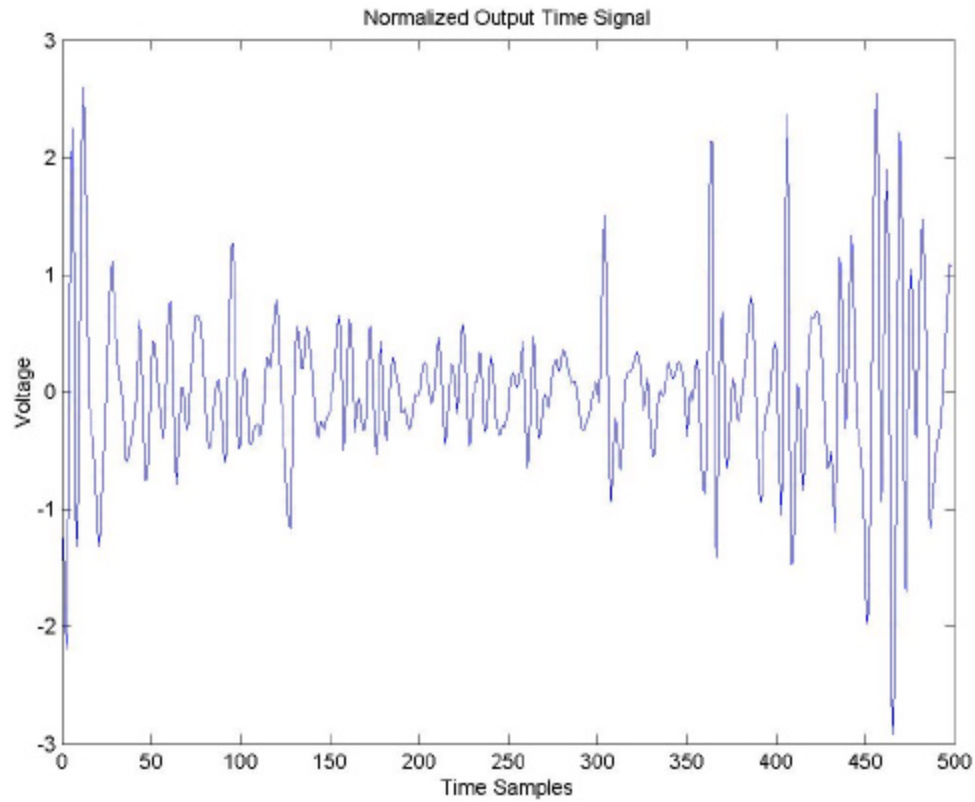
Figure 3.18 Normalized Output

3.3.6 Spectral Estimation

In anticipation of performing Hanning weighting, spectral estimation is performed to enhance the sampled data set of the incoming pulse. Not extending the length of the available pulse data could result in the Hanning weighting reducing some of the spectral components of the target return. Reducing the target return is obviously an adverse effect, so instead the pulse data will be by extended beyond its original sample length. Spectral estimation is performed in this system via the Burg parametric spectral

estimation method [18]. Estimating beyond the existing data window is better than zero-padding, because the range resolution is improved and the full frequency variation is recovered. The result of this operation is an estimated pulse length that is twice as long as the actual incoming pulse data. Figure 3.19 illustrates the effect of Burg spectral estimation on the pulse shown in Figure 3.18.
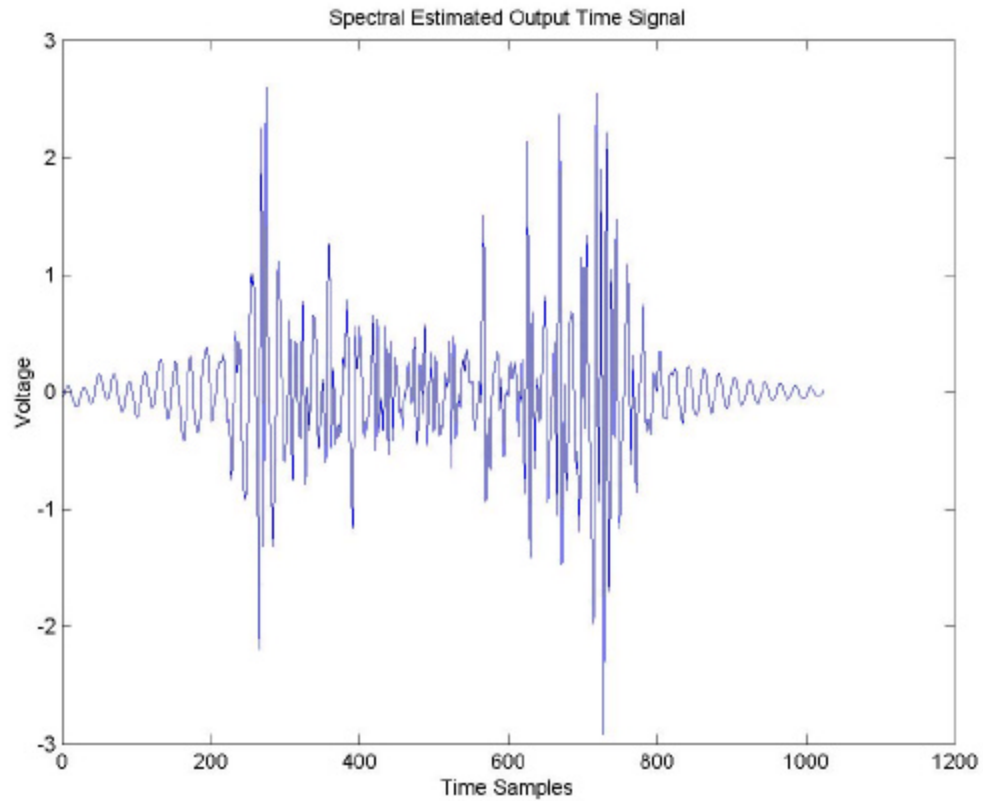
Figure 3.19 Output after Spectral Estimation

3.3.7 Hanning Weighting

Hanning weighting is performed to reduce the sidelobes created from having a pulse of finite length. Without an infinite number of samples, a signal pulse can not be

resolved to its single frequency components. Convolving a sinusoidal waveform with a

rectangle function and finding the Fourier Transform will yield a Sinc function with a

peak sidelobe at –13 dB below the mainlobe. Convolving the same sinusoidal waveform

with a Hanning window reduces the peak sidelobe of the Sinc function to –32 dB below

the peak mainlobe. Hanning weighting reduces the temporal smearing of the echo from a

target by suppressing the sidelobes and improves the resolution of that target. Although

the signal sidelobes are reduced, the Hanning window spreads the frequency response of

the mainlobe. However, there is not a one-to-one correspondence between the sidelobe

suppression and the mainlobe spreading and the result is a reduction of the impact strong

sidelobes have on smaller targets. A Hanning vector the same length as the incoming

pulse is calculated (shown in Figure 3.20) and applied to the incoming pulse. Figure 3.21

illustrates the effect of Hanning weighting performed on the pulse shown in Figure 3.19.
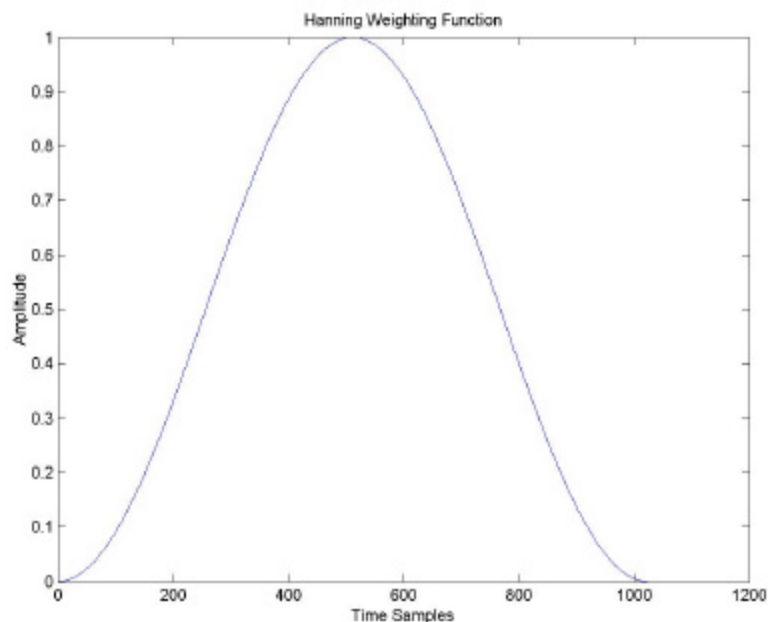
Figure 3.20 Hanning Weighting Function
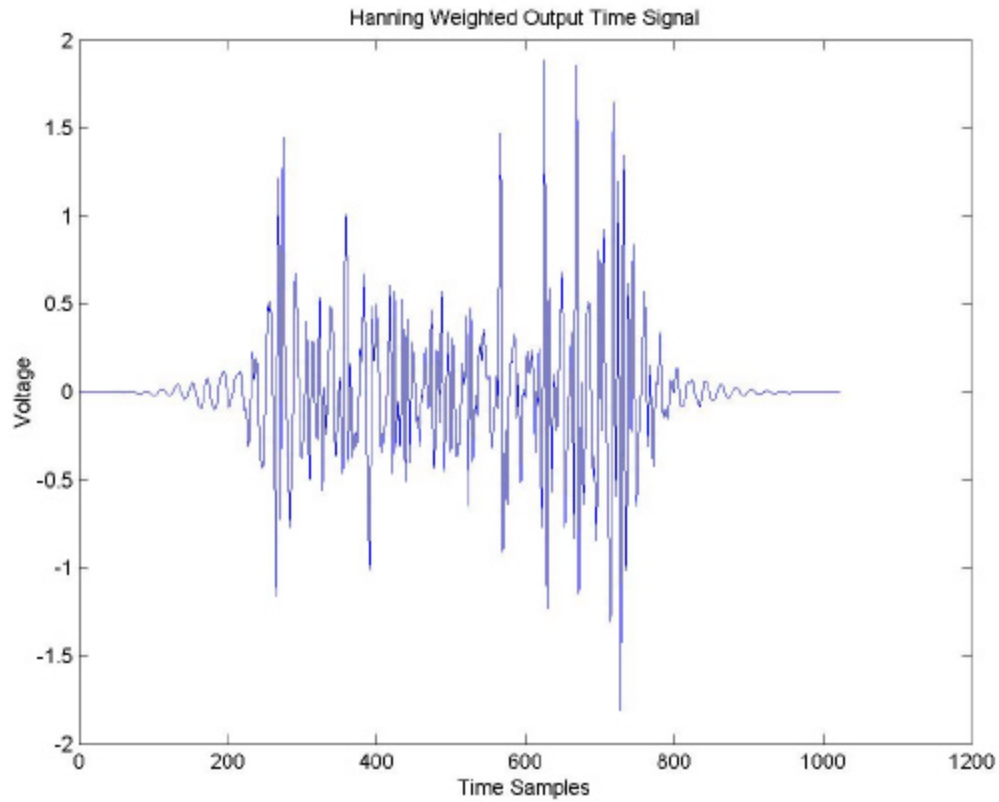
Figure 3.21 Hanning Weighted Output

The demodulated data can be further processed so that the phase information can be used to measure time and thus range. Corrective measures to account for the negative effects of dispersion and ground attenuation have been employed. The direct path signal has been eliminated and the side-lobe peak suppressed to 32 dB below the main-lobe. These processing methods have all been used in anticipation of using the specific SAR processing technique described next. The pre-processed data sent to the SAR processor is a 2D matrix which contains one sampled pulse per receiver location (representing the superposition of all of the reflections received at that location).

36

3.4 SAR Processing

The goal of SAR processing is to calculate the reflection from each individual voxel, as measured by the receiver array. It does this by shifting the phase of the signal from each receiver an appropriate amount so that it is coherent with the data from the other receiver locations. The phase shifts are all based on a common reference point, the transmitter location in this case. The data needs to be reorganized from data organized in a 2D matrix by receiver location into data that can be attributed to specific 3D scatterer locations. The pulse observed by a receiver is the superposition of the reflections of all of the scatterers in the scene contributing to that specific receiver location. Therefore, following the pre-processing described in section 3.3, the return pulse at a given receiver is given by:

$$r_m(t) = \sum_{n=1}^{N} a(t) \cdot \mathrm{Re}[e^{j\boldsymbol{q}_{m,n}}] \qquad (3.1)$$

where m is the particular receiver location, n represents a particular scatterer, N is the total number of scatterers that have a reflection contribution to this receiver, and $\boldsymbol{q}$ is the phase of the receiver signal. The phase, $\boldsymbol{q}$, of that signal is made up of the frequency components of all scatterers that contribute a reflection to that receiver. The phase of a particular transmitter-scatterer-receiver path is given by:

$$\boldsymbol{q}_{m,n} = -2\boldsymbol{p}\left(f_c + \boldsymbol{g}t\right)\boldsymbol{t}_{m,n} + \boldsymbol{p}\boldsymbol{g}t_{m,n}^2 \qquad (3.2)$$

where m is the particular receiver location, n is the scatterer location, f_c is the beginning

frequency of the chirp, g is the chirp rate, t is the pulse length ($0 < t < 2$ ms), and t defines

the path length in time. The receiver pulse variables of interest are the intensity of the

return, represented by the magnitude $a(t)$ in equation 3.1, and the t value in the phase of

equation 3.2. The calculation of this t value will now be discussed.

3.4.1 Path Length Measurements

The first step in finding the contribution from a particular scatterer is to calculate

the 3D path that would have been traversed had a reflection from that scatterer occurred.

An illustration for a single scatterer, or voxel under test (VUT), is given in Figure 3.22.

Tx

Rx

R_n

R_m

Voxel Under Test (VUT)
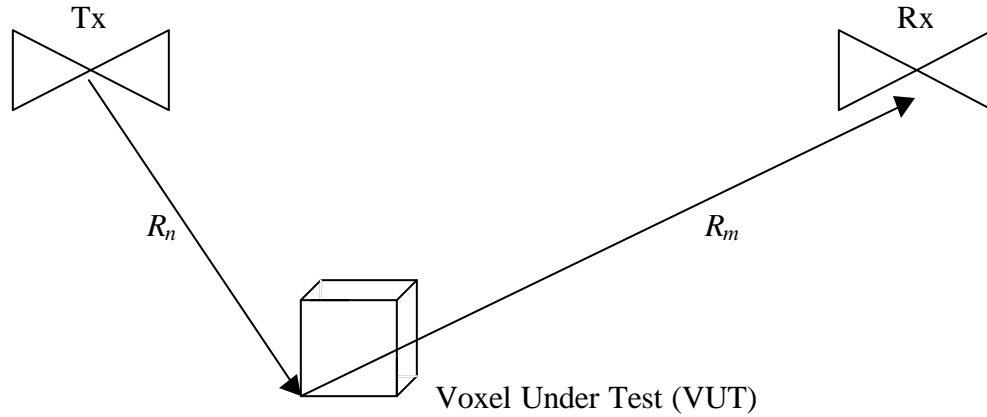
Figure 3.22 Path Length for Scatterer located at VUT location

The variable, t, of equation 3.2 can be described as:

$$t = \frac{(R_n + R_m)}{n} \qquad (3.3)$$

where R_n is the Euclidean distance from the transmitter to the scatterer, R_m is the

Euclidean distance from the scatterer to the receiver, and n is the velocity of propagation

through the medium of transmission. The Euclidean distance is defined as the straight-

line distance between two points and is calculated by:

$$R = \sqrt{\left((X_1 - X_2)^2 + (Y_1 - Y_2)^2 + (Z_1 - Z_2)^2\right)}$$ (3.4)

where X_1, Y_1, Z_1, is the 3D location of the first point and X_2, Y_2, Z_2, is the 3D location of

the second point. The velocity of propagation through a medium is given by:

$$\boldsymbol{n} = \frac{1}{\sqrt{\boldsymbol{m}\boldsymbol{e}}}$$ (3.5)

where $\boldsymbol{m}$ is the magnetic permeability and $\boldsymbol{e}$ is the dielectric permittivity of the medium of

propagation.

Measuring the path length that a pulse travels is a very important step. Since all

of the other phase variables remain constant across every receiver pulse, it is from the

path length information that the specific frequency tone from a particular voxel in the

scene is defined. Voxels in different areas of the scene will have different path lengths

and will therefore have a different frequency tone after mixing with the reference LO

signal. It is in this way that the return of one voxel is differentiated from another.

3.4.2 Phase Correction

Once the path length correction term, $\boldsymbol{t}$, for a particular scatterer-receiver pair is

known, an appropriate phase correction must be found and applied to the receiver data.

The phase correction is employed for a number of reasons. The first is to orient the

signals from the different receiver locations so that they are all coherent. This ensures

that they will add constructively. The second is to determine the reflection contribution

of the current voxel. The phase correction will be different for every receiver location,

depending on the transmitter-VUT-receiver path length. This is because each of the path lengths associated with the receivers is different and all of the other variables mentioned in the above phase calculation are constant for all receivers. This phase correction is applied to produce a signal associated with a particular receiver that is coherent with all other receiver signals for a specific VUT location.

First, the expected phase for a given transmitter-voxel-receiver combination is computed. For a given receiver, m, and scatterer, n, the phase shift is defined by:

$$\boldsymbol{f}_{m,n} = -2\boldsymbol{p}\left(f_c + \boldsymbol{gt}\right)\boldsymbol{t}_{m,n} + \boldsymbol{pgt}^2_{m,n} \tag{3.6}$$

where f_c is beginning frequency of the chirp, $\boldsymbol{g}$ is the chirp rate, t is the pulse length, and $\boldsymbol{t}$ is found by solving equation 3.3 for this particular transmitter-scatterer-receiver path length.

Once the expected phase for the given receiver and voxel is calculated, it must be applied to the input receiver pulse. The phase correction is applied by multiplying the negative expected phase by the phase of the received pulse. This results in:

$$g_{m,n}(t) = \sum_{n=1}^{N} a(t)_{m,n} \, \text{Re}\left[e^{j\boldsymbol{q}_{m,n}} e^{-j\boldsymbol{f}_{m,n}}\right] \tag{3.7}$$

where g is the resultant signal, a is the amplitude, $\boldsymbol{q}$ is the phase of the received signal, and $\boldsymbol{f}$ is the expected phase for a pulse reflecting from the current scatterer.

Since the phase of the received signal, $\boldsymbol{q}$, is composed of the superposition of all of the scatterers that contribute a return at that receiver location, the sum of the products resulting from equation 3.7 produces:

$$g_{m,n}(t) = a(t)_{m,1} \operatorname{Re}[e^{j(\boldsymbol{q}_{m,1} - \boldsymbol{f}_{m,n})}] + \cdots$$

$$+ a(t)_{m,n} \operatorname{Re}[e^{j(\boldsymbol{q}_{m,n} - \boldsymbol{f}_{m,n})}] + \cdots + a(t)_{m,N} \operatorname{Re}[e^{j(\boldsymbol{q}_{m,N} - \boldsymbol{f}_{m,n})}] \qquad (3.8)$$

Whenever the phase of the received signal equals the phase of the current receiver-scatterer (as in the second term of equation 3.8), they will cancel each other out, leaving only the DC amplitude value, *a(t).* When this phase cancellation occurs, it means the current scatterer made a contribution to the signal received at this receiver location. The contribution is contained in the value of the amplitude, *a(t).*

3.4.3 Coherent Summing

The expected phase is calculated for every receiver location with respect to the current scatterer or VUT. Once the phase shift compensation for all transmitter-VUT-receiver paths is performed yielding coherent signals from different receiver locations, the signals are summed together to produce the total signal reflected from the desired VUT location. The processing is repeated for every voxel in the scene. This results in 3D spatial data that is to be processed by the clutter reduction algorithms.

At this point, the collected data has been processed to determine the reflectivity from a particular scatterer in the scene of interest. The clutter in the scene can be estimated and removed, at this stage. The next chapter discusses the theory behind the estimation methods and chapter 5 discusses how they are specifically employed in the clutter reduction process.

Chapter 4

Background Theory

Signal processing algorithms allow data to be conditioned into a form that is conducive to enhancing a particular characteristic of the data. Choice of a time-domain or frequency-domain technique is based on the computation time, complexity, and desired outcome. Prior to exploring the algorithms that are used to reduce the clutter in this GPR system, the relevant theory that provides a foundation to the use of these methods is explained. Reducing the clutter in an image is an important part of GPR processing. Many factors go into choosing a specific clutter-reduction scheme. To a great extent, these factors define the schemes that are ultimately used. For instance, an imposed restriction of this research is not to use *a priori* target or clutter information. After briefly discussing SAR resolution theory, some GPR terminology derived from the demining community will be discussed. Then the theory of time domain and frequency domain clutter reduction techniques employed will be discussed, followed by a discussion of the theory of classical image enhancement techniques.

4.1 Synthetic Aperture Radar

Synthetic Aperture Radar is a particular implementation of conventional pulse radar that is able to generate high resolution results [8, 19]. The effect of SAR on the resolution is apparent along the azimuth dimension. The resolution along the elevation or range dimension is found similar to that of conventional pulsed radars. The range is found by measuring the time difference between transmitted pulse and received pulse,

42

where the bandwidth of the transmitted pulse is proportional to the range resolution. The azimuth direction (which is perpendicular to range) is where the SAR demonstrates its true value.

The resolution in the azimuth dimension of conventional radar is dependent on the beam-width in the azimuth direction. The equation governing the resolution in a conventional antenna along the azimuth direction is given by:

$$r_{con} = \frac{lR}{D} \tag{4.1}$$

where l is the transmitted wavelength , R is the range to the target, and D is the diameter of the antenna along the azimuth dimension. At any substantial range, an impractically large antenna is required to produce fine resolutions. However, the effective aperture size can be increased relatively easily if the antenna is carried on an airborne platform using a much smaller physical antenna. (It should be noted that the same effect is realized with a stationary antenna and a moving target.) Post-processing of coherent waveforms properly aligns the phase from each element and gives the appearance that the data was generated from a single large antenna. The maximum effective length of the synthetic aperture is limited by the width of the target area, however [8, 15]. This length is given by:

$$L_e = \frac{lR}{D} \tag{4.2}$$

where l is the transmitted wavelength , R is the range to the target, and D is the diameter of the antenna along the azimuth dimension. The resolution of the SAR can be found by substituting twice the effective length of the synthetic aperture for the diameter of the

conventional case:

$$r_{sar} = \frac{\lambda R}{2L_e} = \frac{D}{2} \qquad (4.3)$$

where λ is the transmitted wavelength , R is the range to the target, and L_e the effective

length of the synthetic aperture. Therefore, as is shown in equation 4.3, the resolution of

a SAR system is independent of range or wavelength and actually improves as the

diameter of the real antenna is decreased. This is a result of the large beam-width of a

smaller antenna allowing for a larger synthetic aperture [15]. The factor of two emerges

from the two-way propagation path from the antenna element to the target and back in the

synthetic case, as compared to the one-way path of a conventional antenna [8]. This

means that in the SAR case, the signals experience a phase shift both on transmission and

reception such that the resolution is one-half that of the conventional antenna.[2]

4.2 GPR Scans

In an effort to establish a common terminology for describing a GPR system, the

GPR and demining communities developed the terms A-scans, B-scans, and C-scans [1,

2]. These terms will first be defined as they were originally, in terms of a monostatic

GPR system, but then that definition will be extended to the bistatic GPR used in this

system.

In a monostatic system, an A-scan is directly formed from the received return of a

single pulse at a specific antenna location above the ground. The pulse received is the

superposition of the reflections at that receiver location. As the antenna is scanned along

---

[2] Thanks to Larry Adzima for providing a reference, his personal notes, and an explanation comparing
synthetic and real aperture cross range resolutions.

a straight line above the ground, the A-scans are collected and concatenated spatially to form the B-scan. The two-dimensional B-scan can be stacked with other B-scans, taken parallel to the original, to form a three-dimensional data set. The C-scan is a time slice through the B-scan. The C-scan can display the content of a particular B-scan, or all of the B-scans, at a particular range. Figure 4.1 depicts the A-scan, B-scan, and C-scan of a monostatic configuration.
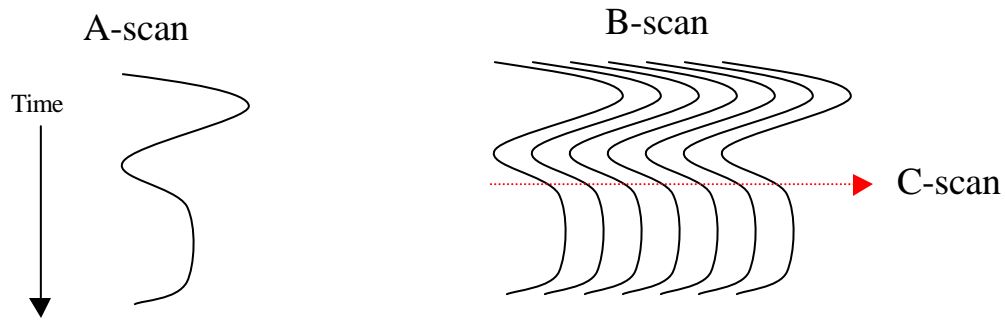
Figure 4.1 A-scan, B-scan, and C-scan in monostatic configuration

The bistatic representation of A-scans, B-scans, and C-scans is analogous to the monostatic representation, with minor adjustments. The A-scan is again a superposition of the reflections from all of the scatterers, received at a single receiver location. However, the B-scan is no longer formed via concatenation of the A-scans. In the bistatic configuration, the B-scan is formed via the SAR processing described in section 3.4. The SAR processing calculates the return from a given voxel. The C-scan is again a time slice through the B-scan. Figure 4.2 illustrates A-scans, B-scans, and C-scans in a bistatic configuration.

45
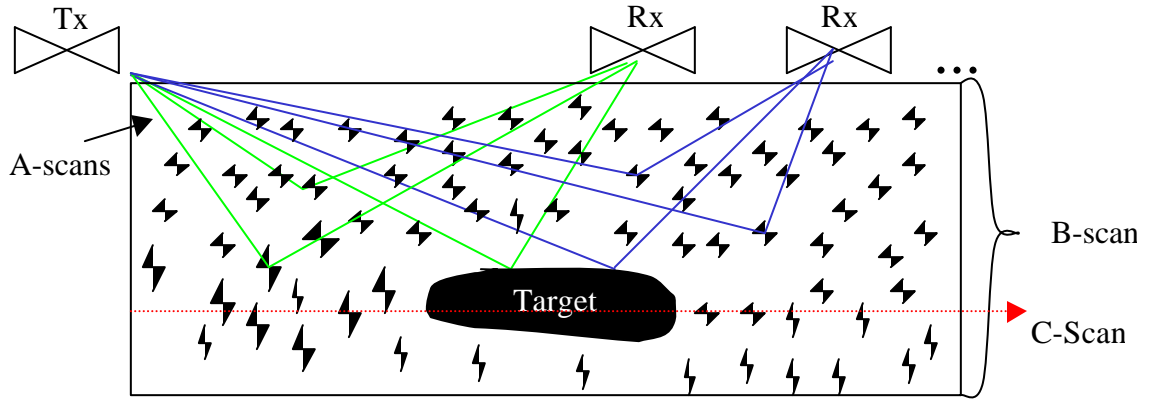
Figure 4.2 A-scan, B-scan, and C-scan in bistatic configuration

4.3 Time Domain Techniques

4.3.1 Mean

In classical probability and statistics, the mean of a vector population can be described as the expected value considering all of the elements in the vector. Extending to the case where *x* can be a multi-dimensional array, recall that the expected value of an array is obtained by taking the expected value of each element [14, 20]. The equation for the mean is described by:

$$m_x = \frac{1}{M}\sum_{k=1}^{M}\mathbf{x_k}$$

(4.4)

where *M* is the total number of elements and *x* is the array of interest.

4.3.2 Root-Mean-Square

The root-mean-square (RMS) operation provides another statistical measure to data sets. The mean given by equation 4.4 doesn't give much statistical insight into data sets that have zero mean. One advantage of the RMS statistic is that it provides insight into data sets that have zero mean. Much like the mean, the RMS of a population of data is referenced in many statistical texts. The RMS of a data array *x* can be described as:

46

$$RMS = \left[\frac{1}{M} \sum_{k=1}^{M} \mathbf{x_k}^2 \right]^{1/2} \qquad (4.5)$$

where the data array x has M elements.

4.3.3 k-Nearest Neighbors

The k-nearest neighbor (kNN) algorithm calculates the density of like voxel values around a given voxel quantity. The number of points to be included, k, is fixed and the volume surrounding the given quantity varies depending on the density of the nearest neighbors. The kNN density estimator can be written as:

$$p(\mathbf{x}) = \frac{k}{NV(\mathbf{x})} \qquad (4.6)$$

where k is the number of points to be included, N is the number of samples, and $V(x)$ is the volume (dependence on $\mathbf{x}$ is explicitly shown).

One application of this algorithm is to use it in decision theory, from a spatial standpoint, using some selected distance measure to measure between an unknown feature vector x and all the training vectors of the various classes in a decision space. If the volume of the space around an unknown feature vector x that contains the k-nearest neighbors from each class is calculated, then x is assigned to the class with the lowest volume (or the class found by the established decision criterion). Once the value for k and a decision measure is chosen, the kNN decision rule can be summed up as follows [21]:

Given an unknown feature vector x,

    1) Identify the k-nearest neighbors from the N training vectors of all classes available.

2) From the k samples found in step1, identify the number, k_i, belonging to each class, where $\sum_i k_i = k$.

3) Assign x to the class with the maximum number of k samples.

4.3.4 Sliding Window

A sliding window has a non-zero impulse response for a finite duration (defined by the size of the window) and can also be called a finite impulse response (FIR) filter. Each output is simply a weighted summation of a finite number of samples of the input sequence [22]. The weighted filter coefficients define the impulse response of the filter. The equation for a causal FIR filter is given by the difference equation:

$$y(n) = \sum_{k=0}^{M-1} b_k x(n-k) \tag{4.7}$$

where $y(n)$ is the present output, b_k is the filter coefficient weight applied at the k^{th} input, $x(n$-$k)$ is the current input, and M is the selected length of the filter. FIR filters can be either recursive or non-recursive.

4.4 Frequency Domain Techniques

4.4.1 Discrete Fourier Transform

The discrete Fourier Transform is a mathematical expression which converts a time or spatial signal into its spectral components. The one-dimensional discrete Fourier Transform is given by:

$$F(u) = \sum_{x=0}^{N-1} f(x) \exp\left(\frac{-j2\textbf{\textit{p}}ux}{N}\right) \tag{4.8}$$

where N is the number of samples and F is the Fourier Transform of f. Since the Fourier

Transform is a separable function, this equation can be applied along each dimension of an n-dimensional data matrix to find the n-dimensional Fourier Transform.

4.4.2 Inverse Discrete Fourier Transform

The inverse discrete Fourier Transform converts frequency information into its time or spatial components. The one-dimensional inverse discrete Fourier Transform is given by:

$$f(x) = \frac{1}{N} \sum_{u=0}^{N-1} F(u) \exp\left(\frac{j2\boldsymbol{p}ux}{N} \right) \tag{4.9}$$

where N is the number of samples and f is the inverse Fourier Transform of F. Similar to the forward Fourier Transform, the inverse Fourier Transform is separable and can also be applied along each dimension of an n-dimensional data matrix to find the n-dimensional inverse Fourier Transform.

4.4.3 Wiener Filter

The Wiener filter is typically used to differentiate a known reference from noise. The Wiener filter restoration transfer function is given by:

$$H = \frac{|R(u,v)|^2}{|R(u,v)|^2 + |C(u,v)|^2} \tag{4.10}$$

where R and C are the Fourier Transforms of the reference and clutter, respectively. Once the filter is calculated, application of the transfer function to the input is done by multiplication in the frequency domain,

$$G = HF \tag{4.11}$$

where G, H, and F are the Fourier Transform of the estimated filtered reference, the Wiener filter transfer function, and the original reference, respectively. After applying

the filter via equation 4.11, the processed reference can be recovered by way of the inverse Fourier Transform of equation 4.9.

4.5 Image Enhancing Techniques

4.5.1 Sobel Edge Detection

An edge can be defined as a boundary between two relatively distinct gray-level regions [20]. Detecting an edge digitally usually involves the computation of a local derivative operator. This is because the derivative of a constant value (or in the case of an image, a relatively unchanging area) is of course, zero. When a local derivative operator traverses an image, values close to zero are computed over areas of relatively little change. When the same local operator lies over a region containing an edge, a non-zero value is found, thus detecting the edge. The value found is scaled by the filter coefficients according to the difference between the gray levels on either side of the edge. Figure 4.3 illustrates an edge and its corresponding derivative functions.[3]

There are many derivative operators used to detect edges. Gonzalez and Woods [20] give examples of the Sobel, Roberts, and Prewitt, in addition to others. The Sobel operators are an attractive choice because they are products of multiplication of a one-dimensional differentiation filter, [-1 0 1], and a one-dimensional low pass filter, [1 2 1] (Figure 4.4). This unique combination produces the differencing effect described above and also gives the effect of a low-pass filter, which provides a smoothing effect. This

---

[3] Gonzalez/Woods, Digital Image Processing, 1992, pp.417. Reprinted by permission of Pearson Education, Inc., Upper Saddle River, New Jersey.
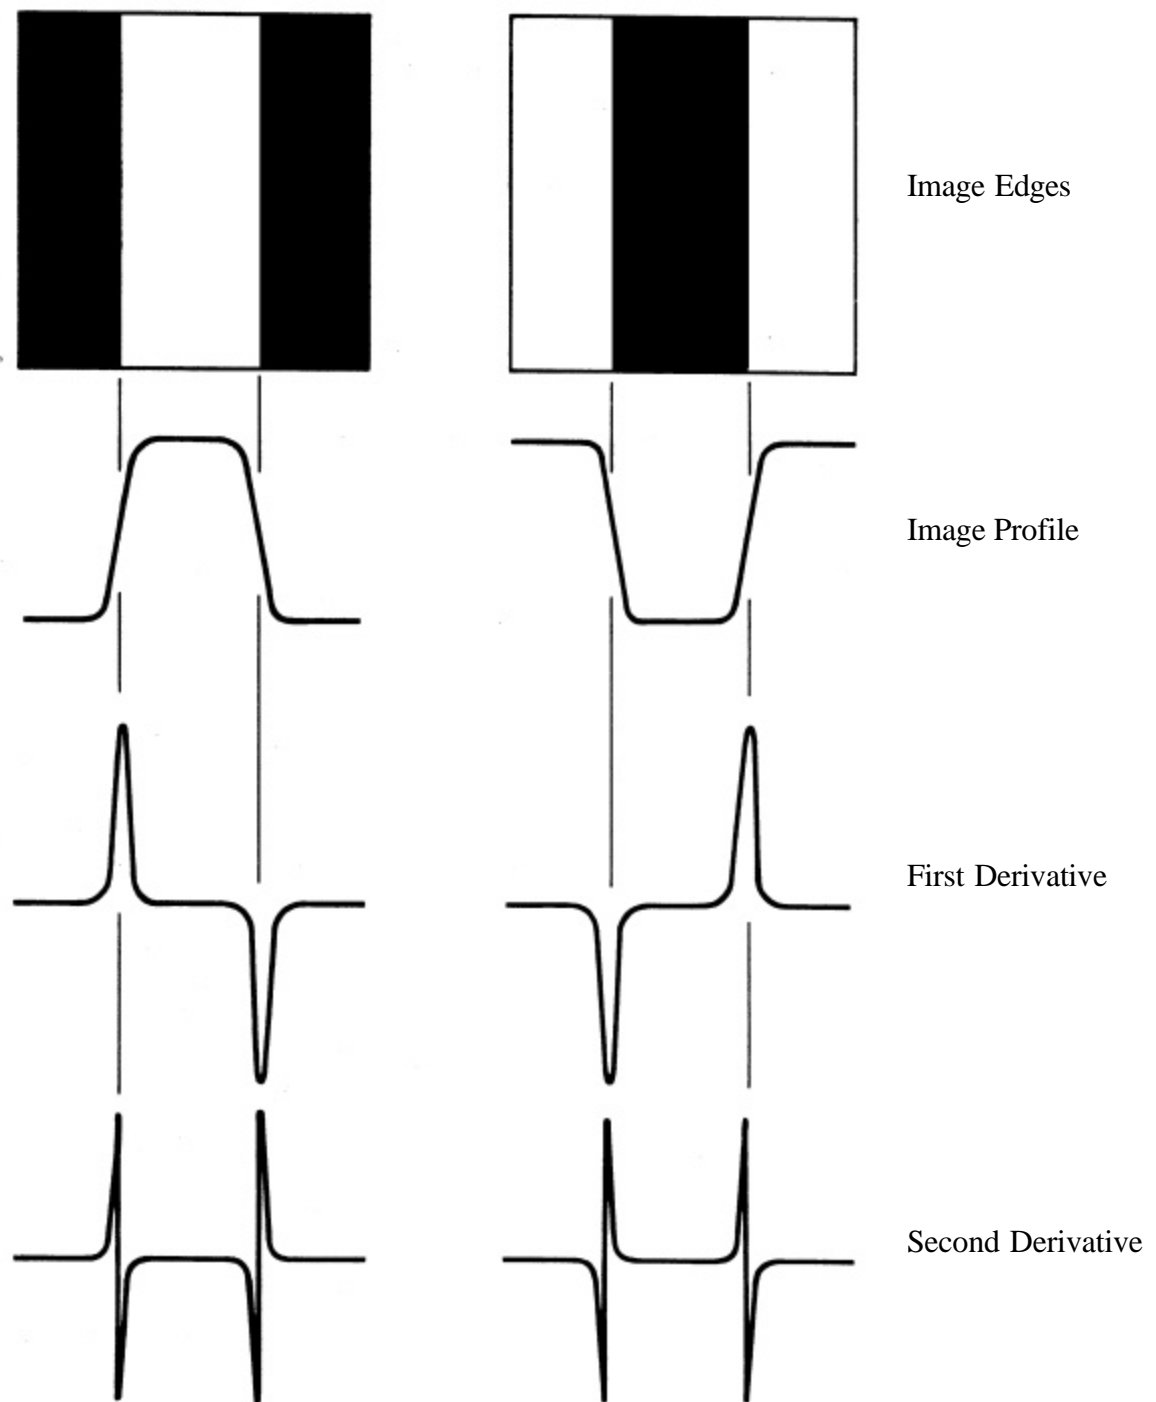
Figure 4.3 Edge Detection by Derivative Operators

smoothing effect reduces the noise that is enhanced by the derivative operation [20]. The

Sobel operators are 3x3 matrices and given in Figure 4.4, with the pixel under test being the center pixel. Together, they provide edge detection in both the horizontal and vertical directions.

$$\begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} * \begin{bmatrix} 1 & 2 & 1 \end{bmatrix} =$$

-1	-2	-1
0	0	0
1	2	1

$$\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} * \begin{bmatrix} -1 & 0 & 1 \end{bmatrix} =$$

-1	0	1
-2	0	2
-1	0	1

Figure 4.4 Sobel Operators

4.5.2 Histogram Equalization

Histogram Equalization is a technique designed to increase the contrast across an image. The way it accomplishes this is to partition all of the pixels of an image into associated gray-level bins (creating the histogram) and then apply a transform to that histogram. The histogram represents the probability density function of a voxel or pixel falling into a certain gray-level value. The goal of histogram equalization is to make the

probability density function uniform. Although a perfectly uniform density is rarely

achieved with real digital images, the histogram equalization transform function results in

a probability density function that approximates a uniform density [20]. This results in

an increase in the dynamic range of the pixels in the image. Histogram equalization can

not achieve a uniform probability density (except in rare circumstances), because a gray-

level bin from the original histogram is not divided into multiple bins of the resulting

histogram. The whole bin is just reassigned to another gray-level bin. The

transformation function in discrete form is given by:

$$T(r_k) = \sum_{j=0}^{k} \frac{n_j}{n} = \sum_{j=0}^{k} p_r(r_j) \qquad \begin{cases} 0 \le r_k \le 1 \\ k = 0,1,...,L-1 \end{cases} \tag{4.12}$$

where T, the transformation function of the bin value r_k, is single-valued and

monotonically increasing, p_r is the probability density function of r, and L is the number

of bins. Once the transformation function given in equation 4.12 is found, the equalized

histogram is a result of applying it to the histogram of the original image.

Now that the theory of the techniques employed against this clutter problem has

been established, the experimental implementation will be discussed next in chapter 5.

Chapter 5

Experimental Implementation

Now that the background theory has been established, the specifics behind the

particular implementation will be explained. The first item for consideration is where in

the image chain would the clutter reduction process be implemented. Initially two

procedures with two separate paths were considered, one that attempted to reduce the

clutter before processing and one that attempted to reduce the clutter after processing.

One of the paths was quickly eliminated because it was impractical in terms of the image

processing employed here. This decision is discussed below. Once it was decided where

the clutter reduction would be applied, a few different techniques to were implemented

reduce the clutter. The specific techniques employed to estimate the clutter were: the

mean, RMS, *kNN*, and an application of the Wiener filter. The code for the techniques

was implemented in MatLab. Each technique is described in detail below. Once the

clutter estimate is calculated, it is filtered from the original input signal. The results of

each of these techniques will be presented and analyzed in the next chapter.

5.1 Injecting Clutter Reduction into the Image Chain

The original consideration was to perform clutter reduction pre- and post-

processing. Experimentation would then take place to determine which method was most

advantageous under a given environment, as it was hypothesized that one method may

not be the most attractive under all circumstances. Based on this premise, the collected

data would be processed via one of the paths in Figure 5.1. It was quickly realized that

performing clutter reduction as pictured in Figure 5.1a is not practical as described

below. Figure 5.1a shows clutter reduction being performed on the raw data before

processing. As a result of the data collection methods and the clutter reduction

techniques employed, there are important processing steps that must be accomplished

before clutter reduction in this GPR SAR system takes place. The raw data is organized

by receiver location and has not been spatially registered to any particular location of the

Raw Data	Raw Data
↓	↓
Clutter Reduction	Processing
↓	↓
Processing	Clutter Reduction
↓	↓
Image Formation	Image Formation
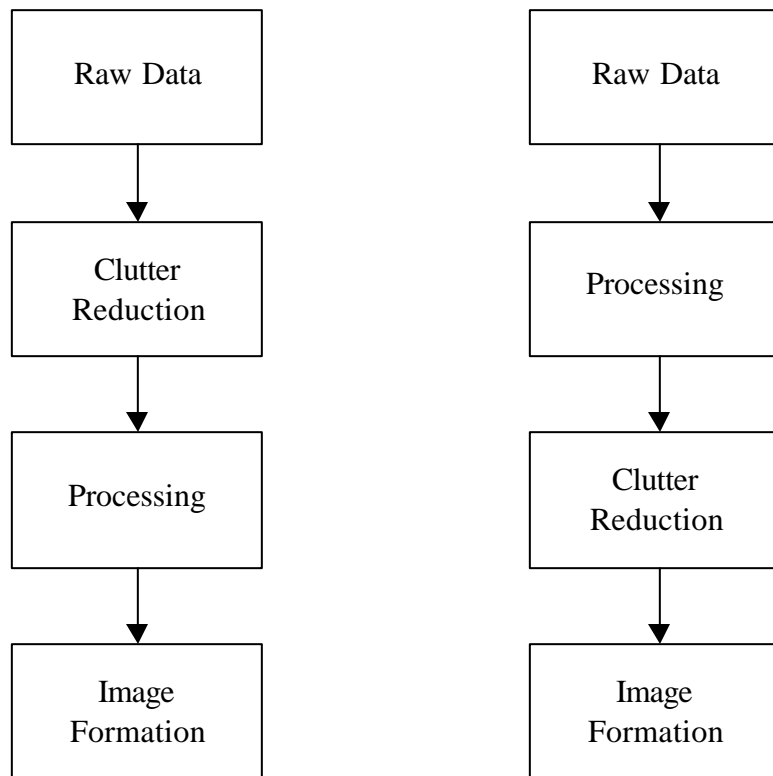
Figure 5.1a (left) and 5.1b (right) Clutter Reduction in the Image Chain

scene before processing. Some of the clutter reduction windowing techniques require

this spatial registration in order to perform properly. The I-Q Demodulation procedure

discussed in section 3.3.2 enables us to measure the phase of the signal collected at the

receiver. This is an important step, because it is this phase information that allows us to determine the length of the path traveled by the transmitted pulse as discussed in section 3.4. Since the path length traveled from the transmitter to a scatterer to a particular receiver location will be dependent on the receiver location, the path length (measured with phase) of a particular receiver pulse will also be dependent on the receiver location. If this change in phase for different receiver locations is not accounted for, false objects will appear in the final image or the reflection will not appear in the true position. The SAR processing procedure discussed in section 3.4 calculates the contribution from a particular scatterer to the pulses of all of the receiver locations. It multiplies each receiver-sampled pulse by an appropriate phase shift, depending on the current VUT location. If any reflection was received from this VUT, it will be indicated there. An alternative to performing clutter reduction operations on the raw received pulses was discussed. It involved performing enough processing to align the phases of the collected data appropriately and then reducing the clutter via one of the methods mentioned. However, in the end it was decided that this would merely duplicate the existing processing and wasn't sensible.

It was discovered that performing clutter reduction operations on data before phase alignment was futile because in order to use the windowing techniques proposed, the data must be spatially registered. Prior to phase alignment, the data is organized by receiver location. If it were not registered, clutter reduction would be performed on a group of real sampled pulse returns which are superimposed on one another. Using a window in this case would lead to inaccurate filter measurements, placing an incorrect

56

value at the VUT. Therefore, the clutter reduction techniques were performed after processing.

When deciding what techniques would be appropriate to address this problem, it was decided only to consider clutter reduction algorithms that do not require *a priori* knowledge of either the target or the clutter. Under ideal circumstances, designing a filter that took advantage of as much *a priori* information as available would be preferred. However, it was determined that the intended use for this system would allow for very little, if any, *a priori* information. Therefore, a restriction to design filters that did not require *a priori* information of the target or clutter environment was imposed.

5.2 Processing Chain Adjustments

The HPF designed to remove the direct path in the pre-processing section of the existing system is discussed in section 3.3.4. It was decided to eliminate this step when constructing an image. The HPF was not included for the following reasons. First, the filter parameters depend on the dielectric permittivity of the specific collection scene. The Rome Laboratories collectors were never able to measure the dielectric permittivity, so the filter parameters specific to a given scene could not be defined. If the parameters are not set correctly, the HPF could attenuate some target data instead of just the direct path signal. Second, it was decided that not including the HPF was advantageous since it was ascertained that including data in the scene that may be direct path was better than risking eliminating target data. If the direct path is included, it can be considered as another clutter signal that needs to be suppressed.

5.3 Estimation-Subtraction Methods

All of the estimation-subtraction methods were implemented two ways. Each method was evaluated using an entire scene and local FIR filter.

5.3.1 Mean

The first method attempted was to simply take the mean of a sample of the scene and subtract it from the VUT. Equation 4.4 was implemented in two ways to accomplish this.

In the first implementation, the sample set is limited to the entire scene. That is, the value of every voxel in the scene is used to determine the mean. Then that mean is subtracted from every voxel in the scene. This operates under the assumption that the clutter is relatively homogeneous across the entire scene. If it is, it will give a similar reflectivity across the entire scene. It is hypothesized that in most cases the clutter signal will dominate the target signal in the scene. If the clutter statistics are such that the mean is representative of the clutter and they are different from the target statistics, removing the mean would eliminate most of the clutter.

The second implementation applies a local FIR filter. The mean of a local sliding window is calculated before subtracting that mean from the VUT. This limits the estimation of the mean to local windows. In this case, the homogeneity assumption of the clutter does not have to apply to the entire seen and can be limited to the local windows. The specifics of this implementation are discussed in section 5.3.4.

5.3.2 RMS

The second method took advantage of another statistical measure, the root-mean-

square (RMS) given in equation 4.5. This method was evaluated because RMS is one of the most common methods to refer to an oscillating waveform. It was also realized that in data sets where the means were close to zero this method would provide some measure of the statistics of the data.

The first implementation limits the sample set to every voxel in the entire scene. Therefore, the RMS of the entire scene is found prior to subtracting it from every voxel in the scene. If the clutter statistics show that the RMS is a representative value for the clutter, then subtracting the RMS of the clutter from every voxel would eliminate most of the clutter.

The second implementation applies an FIR filter to local windows of the data. The RMS of a local window is calculated before subtracting it from the VUT. The specifics of this implementation are discussed in section 5.3.4.

5.3.3 kNN

The third method utilized the *k*-nearest neighbor algorithm described in section 4.3.3, with one major difference. The difference in the implementation proposed here is that it uses the *k*-neighbors nearest to the VUT in value, as opposed to the *k*-nearest neighbors in spatial distance. This method may prove to be useful in circumstances where the clutter demonstrates similar statistics in groups that are physically separated from one another. This application of the *k*-nearest neighbor algorithm is not linear, because the filter taps change as the VUT changes.

The first implementation limits the sample set to every voxel in the scene. The first consideration is to carefully define the number *k*. A major factor to consider when

choosing a *k*-number is the target size. If the *k*-number is not larger than the number of voxels the target occupies then the target could potentially cancel itself out, eliminating the target signal. Obviously, this is undesirable. Once the *k*-number is established, the mean of those *k*-voxels nearest to the VUT in value is found. This mean is then subtracted from the VUT. The *kNN* algorithm would be successful at reducing the clutter assuming the clutter's response had pockets of similarity that were separated spatially across the scene and different than the target. If the *k*-number is chosen appropriately and the *kNN* algorithm successfully represented the clutter statistics, spatially separated clutter would cancel itself out. The target would also be affected, but if the target response to the electromagnetic energy is significantly different than the clutter, the target would still appear after clutter reduction. In other words, the variance within the clutter signal would be less than the variance between the clutter signal and the target signal. An example of this algorithm is given in Figure 5.2. For simplicity, a three-voxel by three-voxel scene is shown. Only the result of the first two voxels will be shown. The voxel circled in red is the VUT and the voxels circled in blue are the nearest neighbors. The mean of these nearest neighbors are found, before subtracting that mean from the VUT.
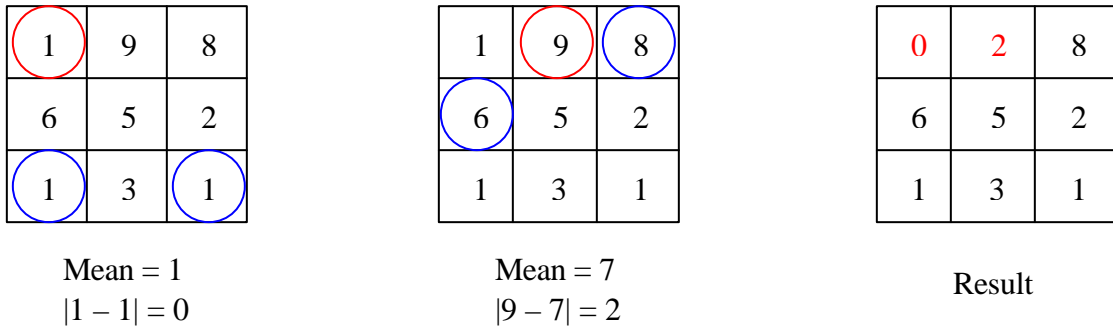
Figure 5.2 *kNN* Algorithm Example

The second implementation applies an FIR filter to local windows of the data. The *k*-nearest neighbors of a local window are found before subtracting their mean from the VUT. The details of this implementation are discussed in the next section.

5.3.4 FIR Sliding Window Filter

In each of the estimation-subtraction cases above, the entire scene is used as the sample set. This operates under the assumption that the clutter environment is homogeneous throughout the scene. Another possibility is that the clutter is homogeneous in smaller regions of data. Performing the estimation-subtraction filters over local areas of data will more accurately calculate the statistics of the data located in close proximity to the VUT. However, the window sizes must be selected carefully. When over the target, the windows need to include some of the clutter in the statistics. If the window is such that this is not the case, the filter will only reflect the statistics of the target and the target will effectively cancel itself out.

In geology, the Law of Original Horizontality states that layers of sediments are deposited horizontally, or nearly horizontally [23]. Since these deposits are potential sources of clutter, they should be a consideration when designing the size of the windows. The windows are designed to stretch across the horizontal deposits in order to get a good sampling of similar sediment. The sizes of the windows are limited to two-dimensional and three-dimensional structures, in an effort to take advantage of the horizontal nature of the earth sediment. According to the Law of Original Horizontality, a horizontal one-dimensional window would most likely only extend across one type of sediment. This would be useful in eliminating that source of clutter. However, the

targets also lie in a horizontal posture. As a result, a one-dimensional window would suppress both target and clutter and potentially eliminate everything in the scene. Hence, one-dimensional windows were not used.

An FIR filter is implemented in each of the cases above: the mean (section 5.3.1), RMS (section 5.3.2), and *kNN* (section 5.3.3). In all of the cases, each of the coefficients of the FIR filter is either zero or one. All of the voxels that are selected to contribute to the specific statistical measure are given a weight of one. A guard window can also be selected to prevent certain voxels in the window from being included in the statistical calculation. The VUT is always in the guard window, but the guard window is not limited to the VUT. All voxels in the guard window have a weight of zero, which does not allow them to contribute to the statistical measure. All of the FIR filters are implemented non-recursively, because the output is not dependent on feedback from previous output values. The VUT is generally located in the center of the window, but is not limited to that location.

5.4 Wiener Filter

The final clutter reduction algorithm utilized in this GPR system was an application of the Wiener filter. This method was implemented in an effort to take advantage of the frequency information contained in the data. Equation 4.10 is used to find the transfer function of the filter and 4.11 to apply it to the input data.

Good performance of the Wiener filter strongly depends on *a priori* knowledge of the exact clutter power spectrum [13]. However, the clutter reduction algorithms proposed here must work under the restriction that information be determined only from

the collected data. This application of the Wiener filter uses the entire input scene (that contains both target and clutter data) to define the filter clutter term. Although this method includes the target signal in the clutter component of the filter, it eliminates the need for *a priori* clutter spectral information. Clutter is assumed to be the prominent source of echo return, so that it will dominate the target return and limit the targets effect in the clutter filter term. This only solves one of the problems in generating the Wiener filter transfer function coefficients without requiring *a priori* information. The generation of a target reference that does not require *a priori* information is still required.

The reference for the target is generated by a diffuse uniform point source reflector model. This model was created as the target simulator designed for this GPR system, described in chapter 3. The inputs to the model include information of the expected size and coarse shape of the target. This is not seen as a violation of not requiring *a priori* knowledge of the target or clutter, because the general size and shape of the expected target would have to be known in order to adequately identify it in the image. Nevertheless, edge detection is being used to define the size and shape of potential targets. The simulator models reflection from all scatterers as completely diffuse. In this way, every receiver location contains a frequency component from every scatterer. Although in a real target environment the receivers in the specular angle will get the strongest return, the returns will not be limited to the specular angle. It is speculated that in a real target environment undulations exist on the surface of the ground-tunnel interface, which will disperse the impinging electromagnetic energy in many directions. The combination of all of the man-made imperfections results in more

diffuse scattering by the target than a perfectly flat target.

The simulator calculates the path length between the transmitter, a single point scatterer, and each individual receiver location using the Euclidean distance measure. The path length measurements are used to define the where the energy came from. Therefore, every receiver location contains a frequency component representing every point scatterer. The phase of the signal is calculated for every transmitter-scatterer-receiver path length according to equation 3.2. Therefore, a given receiver's simulated signal is made up of the superposition of the phases corresponding to the appropriate path lengths of all of the point scatterers in the scene. The same processing described in sections 3.3 and 3.4 is then used to form the image.

The primary targets for this system are underground structures. Therefore, extended targets had to be created from the distribution of uniform point scatterers. Placing the point scatterers at close intervals creates the extended targets. An example of a simulated extended target is given in Figure 5.3. The yellow-green image is the return formed by the processing. The rectangular box is drawn in to show the actual location of the simulated target.

SAR Data 3D Visualization

Image of Simulated Return

Actual Target Location

Z Dimension - Depth (ft)
20
45
70
95
295
245
195
145
95
45
95
45
X Dimension (ft)
Y Dimension (ft)
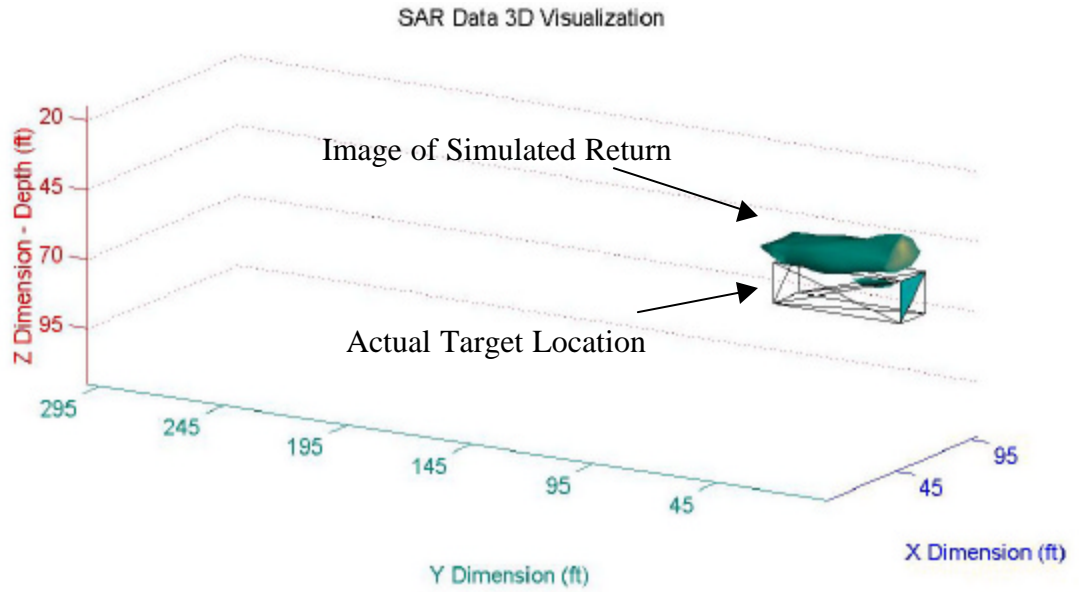
Figure 5.3 Synthetic Extended Target Example

5.5 *Sobel Edge Detection*

Sobel edge detection is utilized as a means to determine possible target sizes and approximate locations since *a priori* information has been prohibited. The edges are found and the size and shape of possible targets are identified. This information is used to determine algorithm parameters such as the size of the sliding window, the *k*-number in the *kNN* technique, as well as the size and shape of the target reference for the Wiener filter. Since the desired targets for this system are underground structures [16], defining edge patterns to identify target structures would be useful. The Sobel operators given in Figure 4.4 are used to detect the edges of possible targets.

The Sobel operators are applied across 2D planes of the image, in both horizontal and vertical directions. A threshold is established to define the minimum magnitude considered to be an edge. If the threshold is exceeded and an edge is detected, a "1" is

placed at the VUT, otherwise a "0" is placed at the VUT. The edge map will display the VUTs that have been labeled with a "1." To illustrate the effects of the Sobel operators, an example using them is illustrated in Figure 5.4.
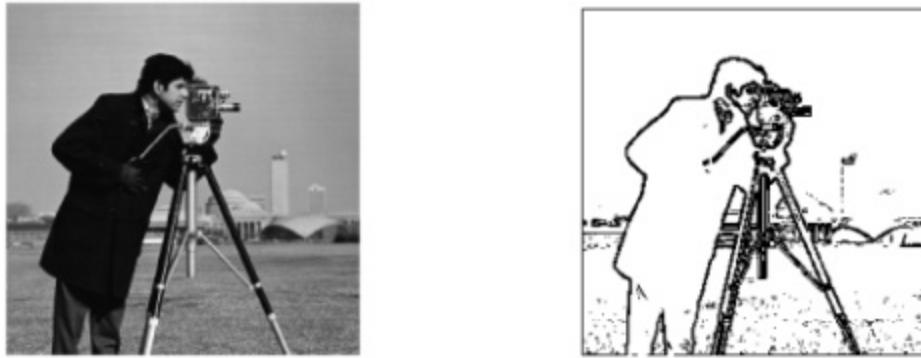
Figure 5.4 Original Image (Left) and Corresponding Edge Map (Right)

<u>*5.6 Histogram Equalization*</u>

Histogram equalization is integrated to improve the contrast for the Sobel edge detection operation. Improving the contrast across an image with low contrast can give the edges of an image more definition. Histogram equalization is performed on the entire 3D scene of the radar images. Equation 4.12 was implemented to find the transformation function required to equalize low contrast images, using 256 bins. To illustrate the effects of histogram equalization, the technique used in this work is applied to an optical image which is presented in Figure 5.5 and 5.6. Notice in Figure 5.5 how the details of the cameraman's jacket are dramatically improved and the sky and buildings in the background have become lighter in the equalized image. These effects can be attributed to the equalized histogram stretching further across the intensity bins than the original

histogram, shown in Figure 5.6.

Figure 5.5   Original Image (Left) and Histogram Equalized Image (Right)

Figure 5.6 Histograms of Images in Figure 5.5 above

Chapter 6 gives the results of the algorithms discussed in this chapter applied to

data sets collected with the GPR system described in chapter 3. Chapter 7 then discusses

the implications of these results and presents some conclusions.

Chapter 6

Results

It was decided to measure the results in a quantitative and qualitative way. Both quantitative and qualitative metrics have a contribution when analyzing the performance of the image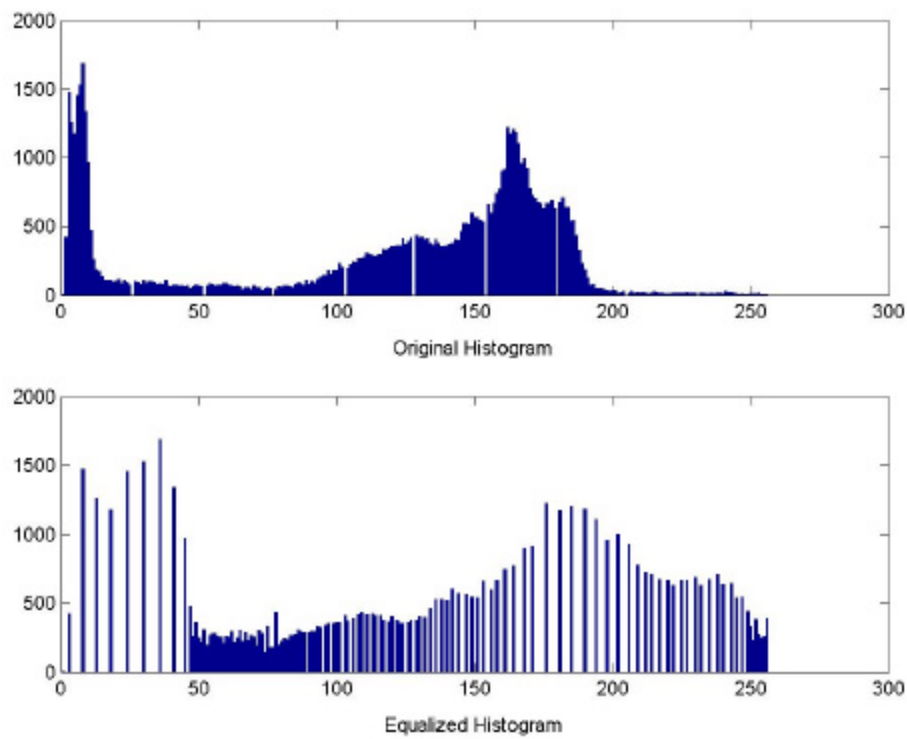 processing schemes presented here. A quantitative metric was chosen in order be able to measure the performance of the clutter reduction filters mathematically. This is important, because it provides quantitative metrics for comparison. However it should be noted that, although somewhat subjective, viewing the final image is also an important criterion to use when analyzing the results of image processing techniques. That is why the final images themselves are also included in this section. It is believed that these quantitative and qualitative measures complement each other well.

The quantitative metric that was decided on first defines an area of voxels around the tunnels as the target signal. The clutter signal was defined as the remaining voxels in the scene. The ratio of the target signal magnitude and the clutter signal magnitude before and after each clutter reduction method will be compared. The ratio is expressed by:

$$I = \frac{|t(x, y, z)|}{|c(x, y, z)|} \tag{7.1}$$

where $t(x,y,z)$ is the magnitude of the three-dimensional target and $c(x,y,z)$ is the magnitude of the clutter. This will measure the improvement in target signal magnitude with respect to the reduction of the clutter.

6.1 Collection Area Descriptions

6.1.1 Nevada Test Site

The United States Department of Energy (DOE) established the Nevada Test Site (NTS) in the 1950s as an on-continent proving ground for nuclear weapons testing. That testing continued for more than four decades until the moratorium on nuclear weapons testing in 1992, when the DOE diversified the use of the NTS. Among many other governmental systems tested at the NTS, this GPR system was recently tested there. The actual site at the NTS used for data collection was the U16b Tunnel Complex, Portal #3. There is no metal mesh lining on the ceiling of this tunnel, as in commercial mines. Reflection in this tunnel is completely caused by the tunnel-air interface. The ground is primarily made up of fine desert gravel. At the time of collection, the surface was wet from rainwater. Water, with a dielectric permittivity around 80 will definitely affect the penetration depth by attenuating the transmitted signal.

A 100-ft by 300-ft receiver array was used for collection at the NTS tunnel. It was positioned directly over the tunnel, which is buried under a sloping hill. The receiver array was placed on the hill to go along the longest dimension of the tunnel. The tunnel itself is about 16.5 feet wide and 16.5 feet high (Figure 6.1). Due to the size of the receiver array, three transmitter locations had to be used, one for each 100-feet by 100-feet receiver sub-array. Each transmitter was centered approximately 60 feet from the center of the corresponding sub-array. The origin of the 3D scene is set at the first receiver location (bottom-left of Figure 6.2). Figure 6.2 gives the geometry of the tunnel, transmitters, and receiver array.

Figure 6.1 NTS Tunnel Portal
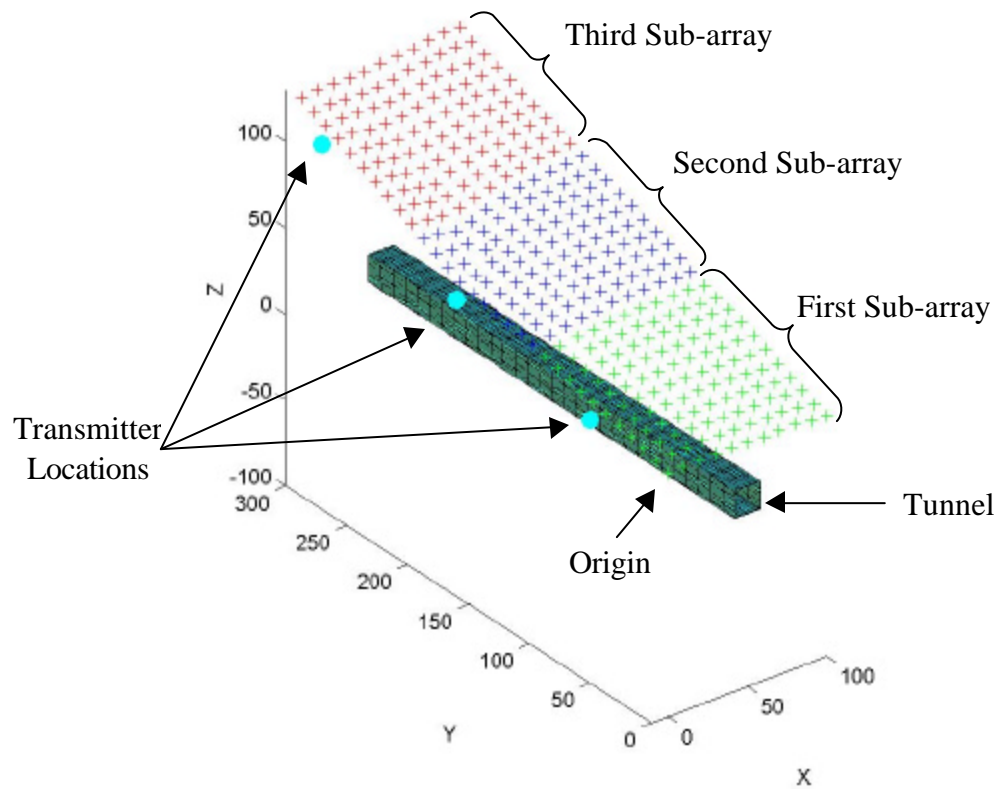
Figure 6.2 NTS Tunnel Orientation

The hill could definitely affect the results, because the tunnel was about 130 feet deep towards the end of the array and only 30 feet deep at the beginning of the array. The signals received at the end of the array will experience much more attenuation than the signals received at the beginning of the array. Combine the hill with the wet conditions and there is a question as to whether the receivers near the end of the array even contributed to the image scene. Therefore in addition to the full array, it was decided to show results from the shallowest sub-array (the green sub-array of Figure 6.2). The tunnel is only 30-65 feet below the receivers of the first sub-array. This was done in order to either support or refute the theory that the attenuation experienced by the data collected further up the hill impeded it from contributing to the image scene. By adding the results from this sub-array, it is hoped that the argument will be elucidated.

Figure 6.3 shows the full-array original data for the NTS tunnel site. Figure 6.4 illustrates the original data collected from just the first sub-array at the NTS. Notice that
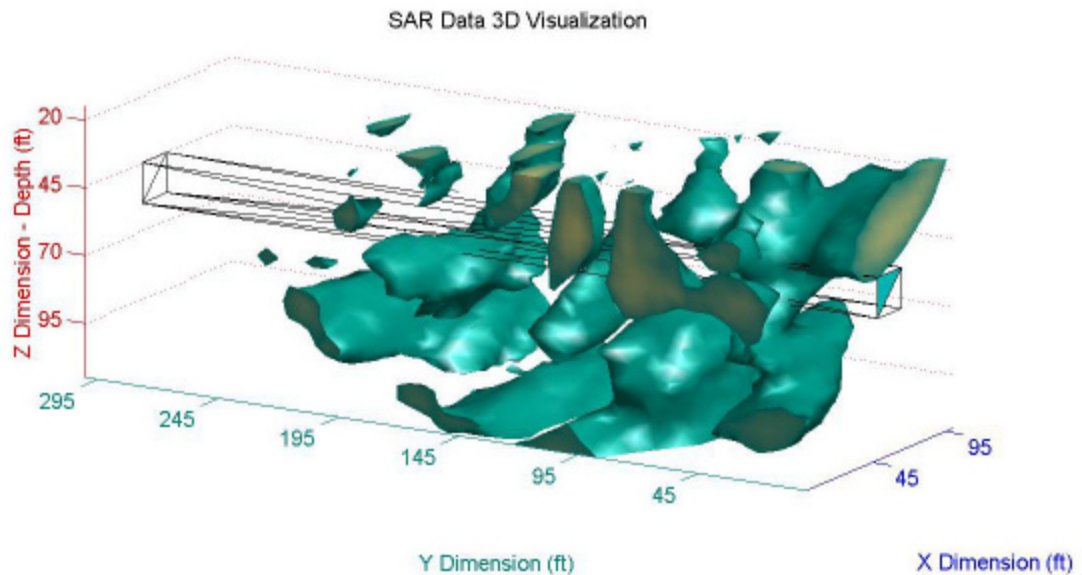
Figure 6.3 NTS Original Scene, 100x300 Array

the tunnel is drawn in the correct location and with the proper dimensions. This will

continue to be illustrated to show the actual tunnel location in all of the images involving

data collected in the field. Applying the proposed metric to the original scene in Figure

6.3 yields a target signal magnitude to clutter signal magnitude ratio of .2155. Out metric

yields a value of .2693 for the sub-array data set in Figure 6.4.

SAR Data 3D Visualization
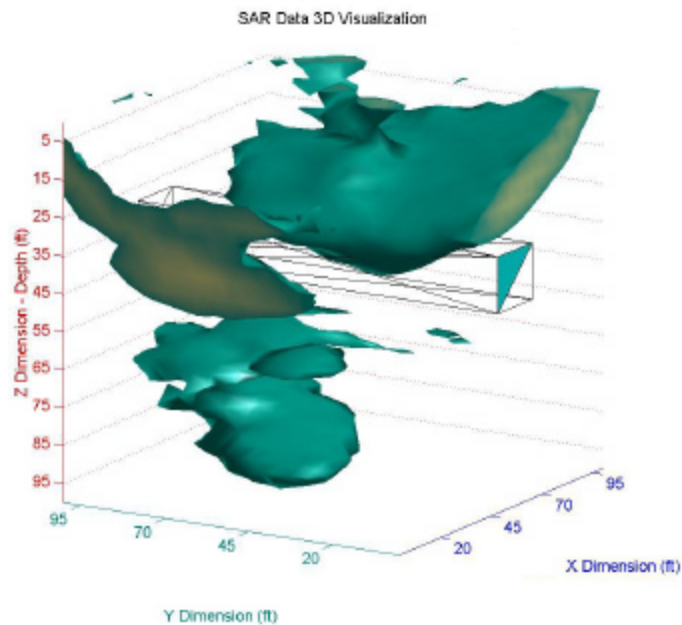
Figure 6.4 NTS Original Scene, 100x100 Array

6.1.2 Zinc Corporation of America

Zinc Corporation of America (ZCA) has a facility located in Balmat, New York.

There are three mines in the Balmat facility: Pierrepont, Hyatt, and the number four

mine. The data presented here was collected at the Hyatt mine. The mineshaft ceilings

are lined with a mesh metal material to keep loose rock from falling on miners and

equipment as they travel to and from the ore bodies. This mesh serendipitously allows

for the potential of improved backscatter reflection, because the metal is a good

conductor. At the time of collection, the Hyatt mine was partially filled with water. The pump that prevented this is the past was turned off, because the mine had not been operational for about nine months before the collect. The water filled the mine tunnels up to the level of a nearby river. The ground environment is made up of marble, quartz, and magnesium carbonate.

A 100-feet x 100-feet receiver array was used at the Hyatt mine. It was positioned on a flat piece of land directly over two mine drifts that were approximately 130-feet and 150-feet deep, respectively (Figure 6.5). The dimensions of the tunnels are

Figure 6.5 ZCA Hyatt Mine Layout

approximately 16 feet high and 16 feet wide. A single transmit antenna was located 65 feet from the center point of the receiver grid. The image of original scene is presented in

Figure 6.6. Although there are two drifts that appear in the scene (Figure 6.5), it was decided to only drawn the drift that is 130-feet deep so that the view isn't unnecessarily obstructed. The deeper drift does not show up as strong as the shallower one, but it does appear in the image as well. Also, notice that the x and y dimensions vary from 0 to -100. The data was processed with these dimensions, because of a suspected sign error. When processed from 0 to 100 in the x and y directions, the data set yields a very small reflection near the surface. There are no other reflections in the remainder of the scene. Through experimentation, the configuration in Figure 6.6 was used to display the data
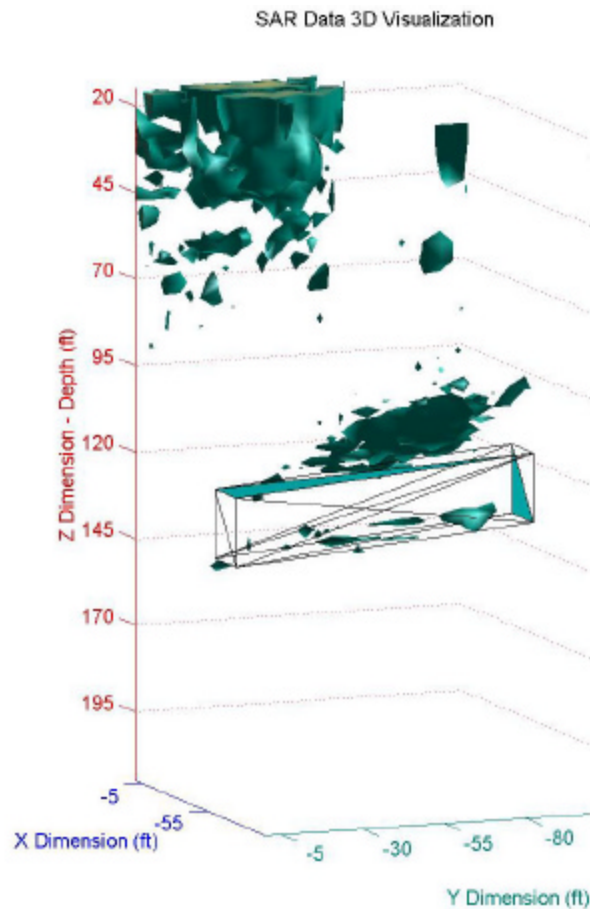

Figure 6.6 Original ZCA Scene

and it was ascertained that the data was originally created with a sign error. So, the dimensions of Figure 6.6 were used. The metric produces a value of .1576 for this scene.

6.2 Image Enhancing Results

6.2.1 Sobel Edge Detection

Some of the clutter reduction methods require inputs regarding the size and shape of the expected target. Sobel edge detection was used so that appropriate parameters required by the algorithms could be defined. Figure 6.7 illustrates the edge detection technique applied to the 100-feet by 300-feet NTS data set with a threshold of 3.5.

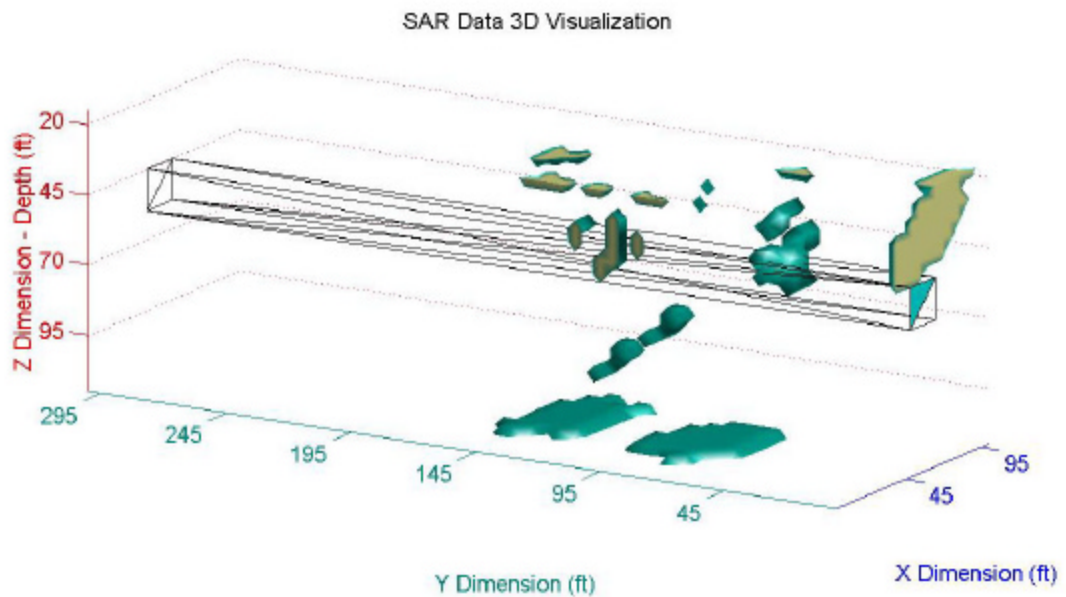**Figure 6.7** Edges of NTS Scene, 100x300 Array

It was determined that there were three possible targets. The first suspected target stretches from 40 to 85 in the x-direction, 0 to 60 in the y-direction, and 10 to 55 in the z-direction. The second target extends from 25 to 70 in the x-direction, 60 to 100 in the y-direction, and 65 to 100 in the z-direction. The third possible target is positioned from 30

to 80 in the x-direction, 120 to 155 in the y-direction, and 0 to 100 in the z-direction. The

first target mentioned is in the immediate vicinity of the actual location of the tunnel,

although is it larger than the tunnel in the x and z dimensions.

Figure 6.8 shows the results of running the edge detection algorithm on the data

from the first sub-array. In this example only one possible target appeared. The

threshold was set to 2.8 in this example. The probable target extends from x coordinates

of 40 to 80, y coordinates of 0 to 70, and z coordinates of 10 to 30. This target is in the

same location as the tunnel, but is larger than the actual tunnel in the x dimension.
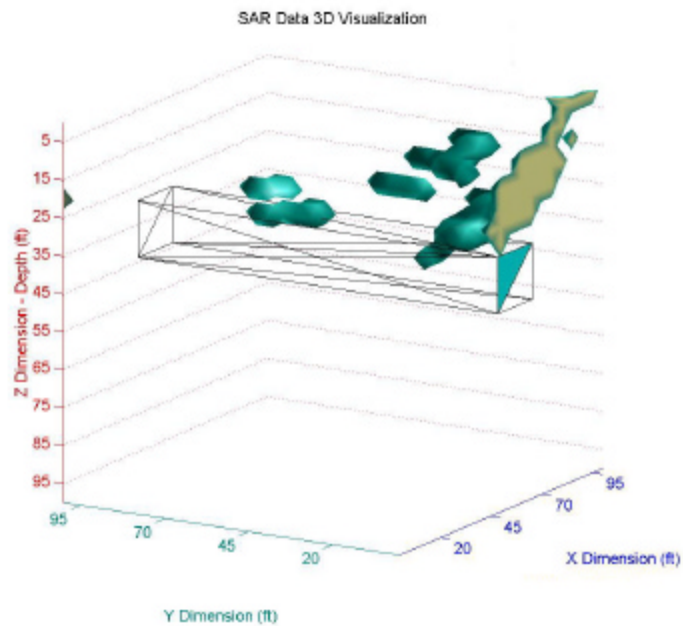

Figure 6.8 Edges of NTS Scene, 100x100 Array

Sobel edge detection produced similar results when applied to the Hyatt mine data

set. The threshold used was 2.0, which was the lowest of any of the data sets. This

should not come as a surprise, because one should expect the attenuation of the pulse

reflecting off of the tunnel to be much greater in this case, due to its depth. The edge

map is shown in Figure 6.9. The first possible target extends from 0 to –35 in the x-

direction, 0 to –50 in the y-direction, and 0 to 50 in the z-direction. This target is very

close to the surface. The second suspected target has dimensions that stretch from –40 to

–80 in the x-direction, –30 to –80 in the y-direction, and 95 to 125 in the z-direction.

This target in is the area immediately above the actual tunnel location.
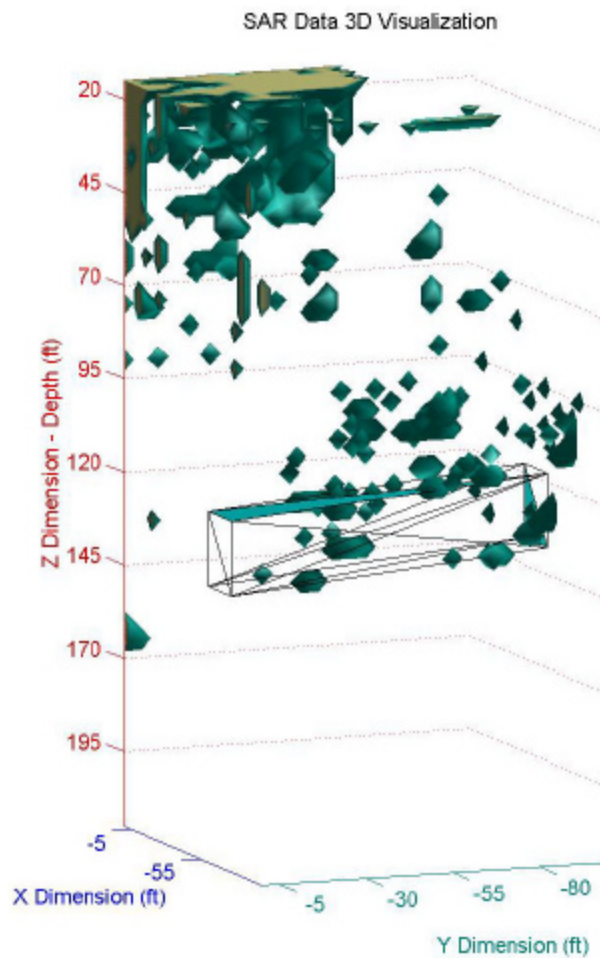
Figure 6.9 Edges of ZCA Scene

Histogram equalization was employed to increase the contrast of the images to improve edge detection. The original and equalized histograms of the NTS 100-feet by 300-feet data array are shown in Figure 6.10. When the equalized histogram is applied to the data, the resulting image is shown in Figure 6.11. Figure 6.12 shows the result of applying the edge detector to the equalized image. The edge detection threshold was normalized to a value of 5.7869 to account for the increase in contrast. Unlike the
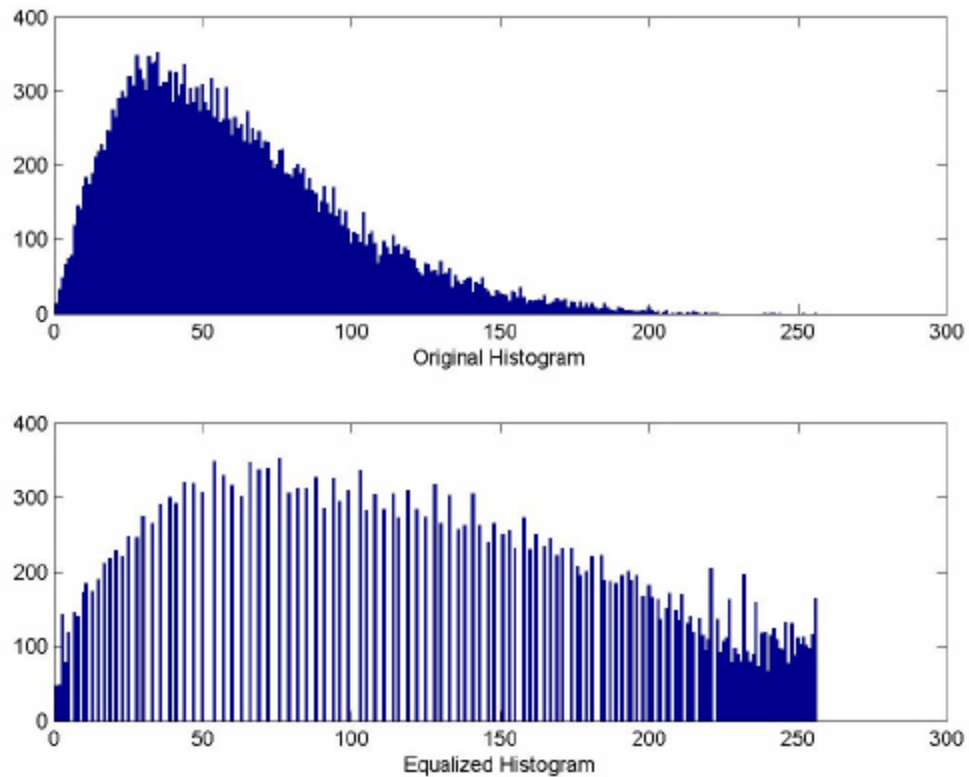
Figure 6.10 Histograms of NTS Scene, 100x300 Array

Figure 6.11 Equalized NTS Scene, 100x300 Array
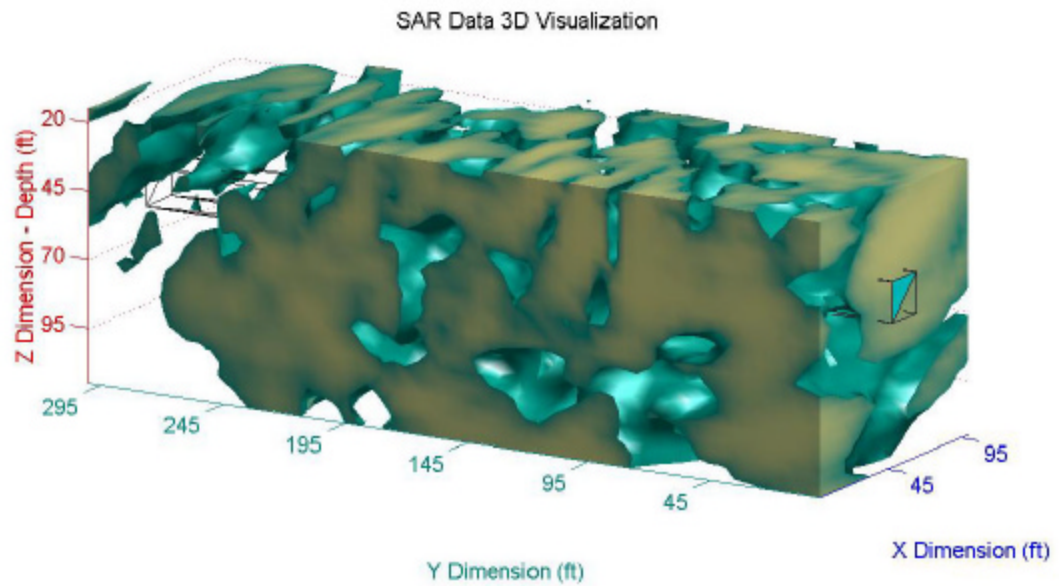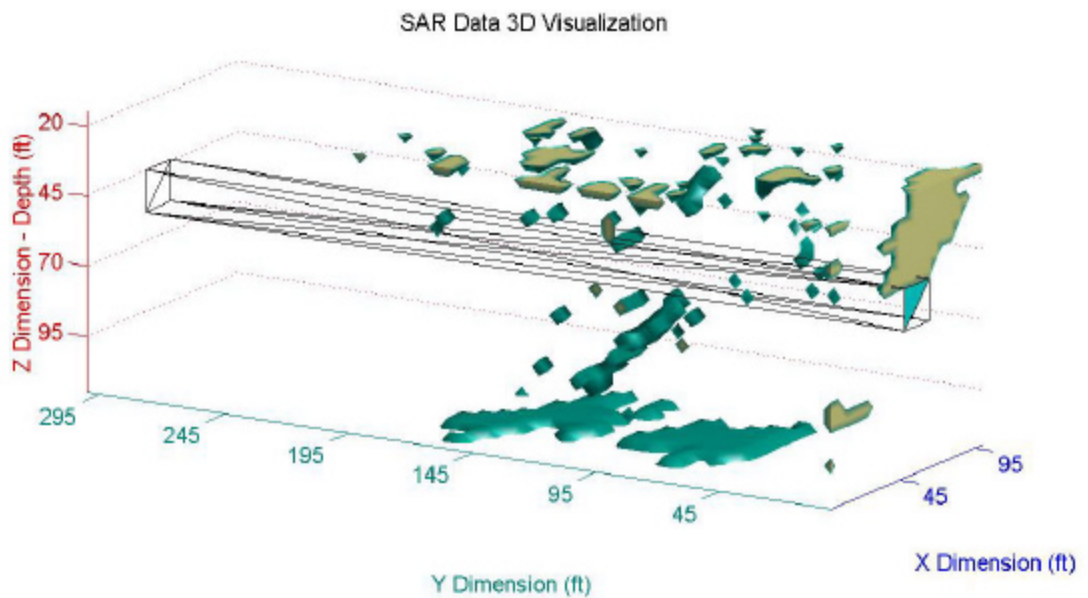

Figure 6.12 Edge Detection following Histogram Equalization, NTS Scene 100x300

cameraman example presented in the last chapter, target detail improvement is not

evident. In fact it may even be obscured. Looking closer at how these images are formed

is needed in order to realize the effect histogram equalization has had on the resulting

image. The SAR processing described in chapter 3 measures the returns from every voxel, given the components of all of the received signals. Every voxel is assigned a value corresponding to the signal magnitude measured by the receivers. By increasing the contrast through histogram equalization, the magnitude value of a given voxel is changed. The histogram of the scene presented has, as a large majority, low magnitude voxels. In fact, only 6.18 % of the voxels have a magnitude level that exceeds the value of the center bin in the original histogram (see "Original Histogram" in Figure 6.10). As a result of performing histogram equalization, many of these voxels are put in bins of increased magnitude. Consequently in the equalized histogram, 50.5 % of the voxels have a magnitude value greater than the center bin. Since the image viewer plots the magnitude of the signal intensity assigned to a particular voxel, the equalized image appears to have many more returns than in the original image. In the cameraman images of chapter 5 the pixel values are grayscale. So by moving bins you're making a pixel either lighter or darker. The same phenomenon is happening in this radar image. In this radar image, signal magnitude is either added or subtracted from a particular voxel. However this can cause false targets to appear, or remove weak targets from the scene. There appeared to be no improvement in detecting targets in the scenes by Sobel edge detection when histogram equalization was employed, however, there may be scenes in which this is useful. This result is consistent with what was seen for the other sets of data; therefore, images of those results are not presented.

6.3 Clutter Reduction Results

6.3.1 Mean

6.3.1.1 NTS with 100x300 Receiver Array

The first experiment was to calculate the mean of each voxel and subtract it from every voxel in the scene. Figure 6.13 shows the mean-subtraction filtered data of the full-array of NTS data. Although a large return is experienced at the beginning of the actual tunnel location, various clutter returns still leave doubt as to a definitive target.
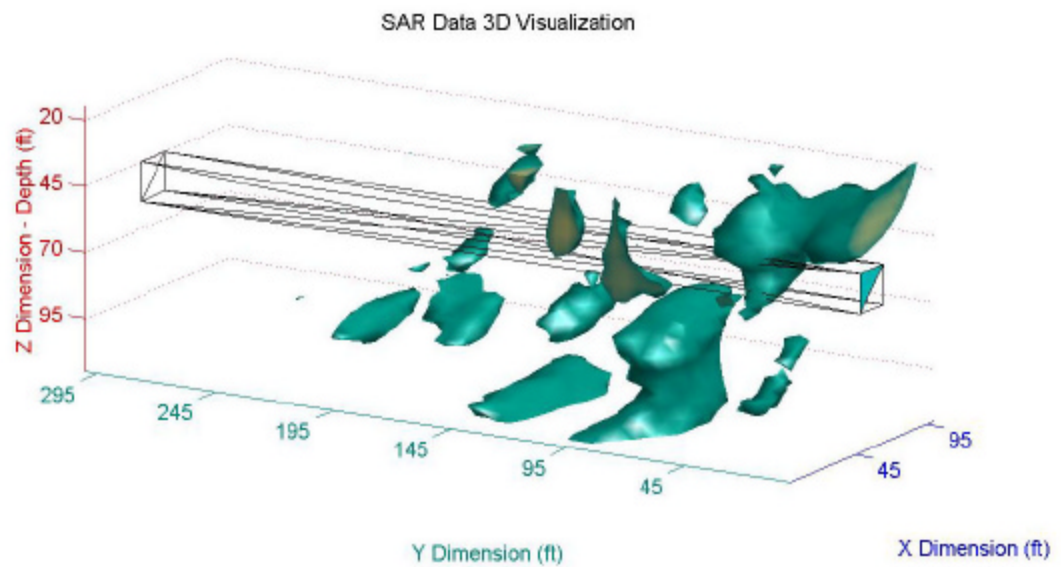

Figure 6.13 Mean Reduced NTS Scene, 100x300 Array

6.3.1.2   NTS with 100x100 Receiver Array

The next experiment performed the mean-subtraction algorithm on the data from the first sub-array of the NTS data. The image resulting from mean-subtraction is displayed in Figure 6.14. The outcome from the sub-array produces a very strong return

from the tunnel and very small patch of clutter along the plane where x is zero. This

image certainly promotes the tunnel location for the target.
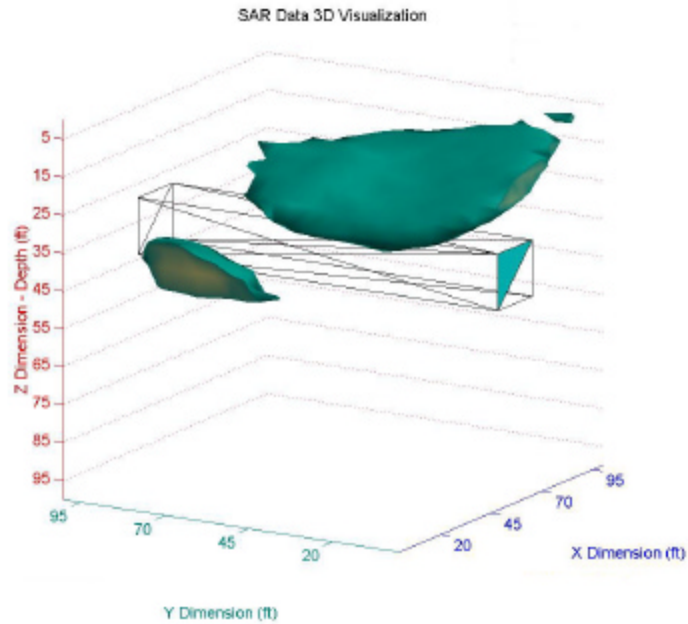
Figure 6.14 Mean Reduced NTS Scene, 100x100 Array

6.3.1.3 ZCA

Subtracting the scene mean from every voxel in the ZCA scene yields the result of

Figure 6.15. Some of the spurious clutter from the original scene has been eliminated,

but the outcome of this method only leaves a small reflection from the actual target

location and has diminished the target quite significantly. There are two definitive

possible targets in the resulting image.
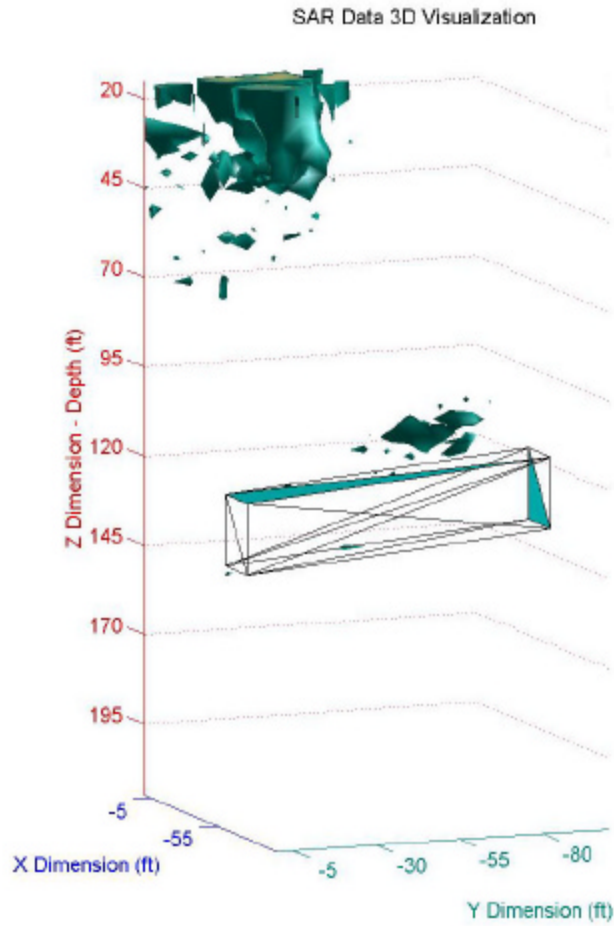
Figure 6.15 Mean Reduced ZCA Scene

6.3.1.4 Mean Window Filters

The windowing technique was applied to all three of the estimation-subtraction algorithms. Through experimentation, it was found that there were similarities between all of the estimation-subtraction window techniques. Therefore, in order to adequately display the results illustrating the effect of applying the sliding windows without adding unnecessary redundancy, a different window will be shown in each estimation-subtraction method. The window chosen for the mean-subtraction technique is two-dimensional with dimensions of 7-voxels in the x-direction, 1-voxel in the y-direction,

83

and 7-voxels in the z-direction. The VUT is the voxel directly in the center of the window. The result of applying this window to the full-array of NTS data is shown in Figure 6.16. There is a large return at the beginning of the actual tunnel location, but

SAR Data 3D Visualization
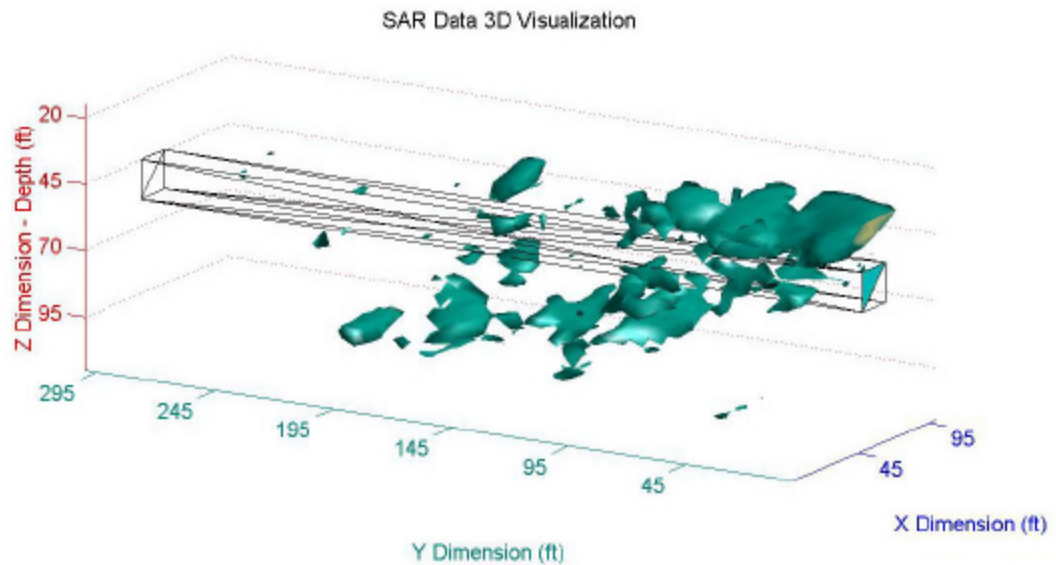
Figure 6.16 Mean Window Reduced NTS Scene, 100x300 Array

quite a bit of clutter remains.

When this window is applied to the first sub-array of the NTS data, the image of Figure 6.17 results. A large return immediately above the actual tunnel location is evident. This image would seem to suggest a single target coincident with the real tunnel location.

The result of applying this window to the ZCA scene is displayed in Figure 6.18. The clutter near the surface in the original image is significantly attenuated, while the reflection appearing at the tunnel location is still very visible. The reflection from the deeper tunnel in the scene is even slightly enhanced by this windowing technique.

Figure 6.17 Mean Window Reduced NTS Scene, 100x100 Array

Figure 6.18 Mean Window Reduced ZCA Scene

Table 6.1 displays the results of the metric for the mean techniques. The clutter magnitude is not reduced as much as the target in the full-array NTS and ZCA scenes for the total mean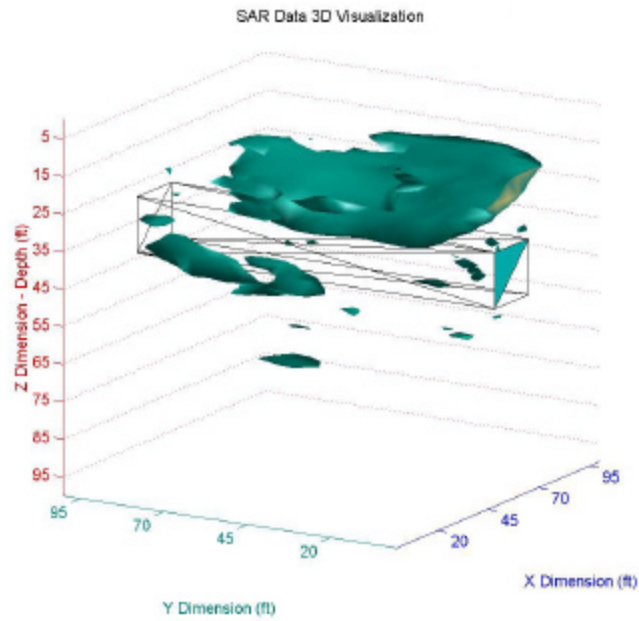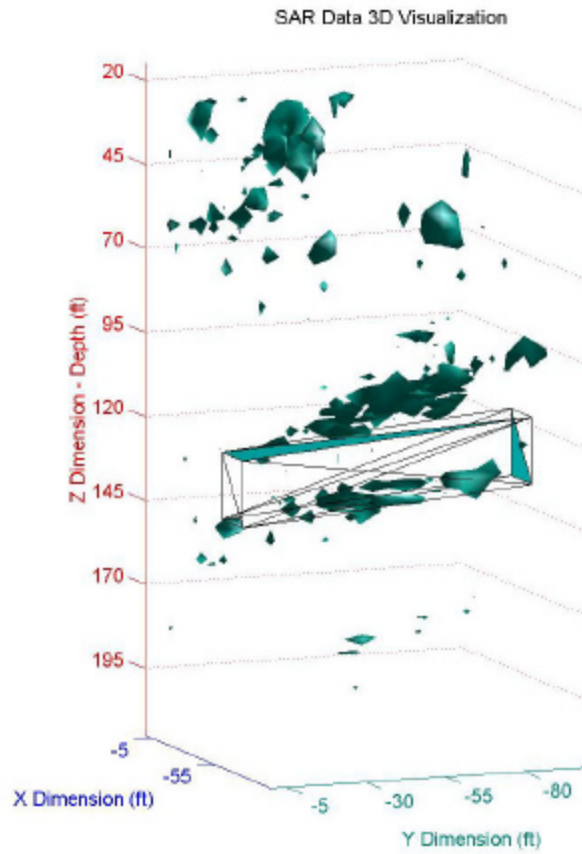 technique. This result is expected for unfiltered scenes where the average target voxel has a magnitude that is below that of the mean of the scene. The total mean technique yields an improvement in the target magnitude to clutter magnitude ratio for the data of the NTS sub-array.

Method	Full-Array of NTS Data	First Sub-Array of NTS Data	ZCA Data
Unfiltered Data	.2155	.2693	.1576
Total Mean	.2083	.3419	.1328
Mean Windows	.4837	.6316	.2630

Table 6.1 Metric Summary for Mean Techniques

In each case where a local window was used the filtered data had a greater target magnitude to clutter magnitude ratio. According to the metric, the technique had the greatest impact on the first sub-array of NTS data. This is reflected by the visual images as well. Notice that the images show that the ZCA mine data also experienced visual enhancement after filtering with the windows. This is expected, because the ZCA mine is the deepest of the three. The energy reflecting off of this tunnel will experience more attenuation than the clutter closer to the surface. The window amplifies the target return (as compared with the peak magnitude voxel), by preventing the statistics of the large clutter near the top of the scene to corrupt it. The window is also small enough to reduce the large clutter at the top of the scene.

The RMS-subtraction method was then performed on the full-array of data from the NTS. The resulting image is presented in Figure 6.19. Clutter is visible, but some of the smaller clutter returns have been reduced from that of the mean-subtraction example. A large return at the beginning of the tunnel is present again as well.

SAR Data 3D Visualization
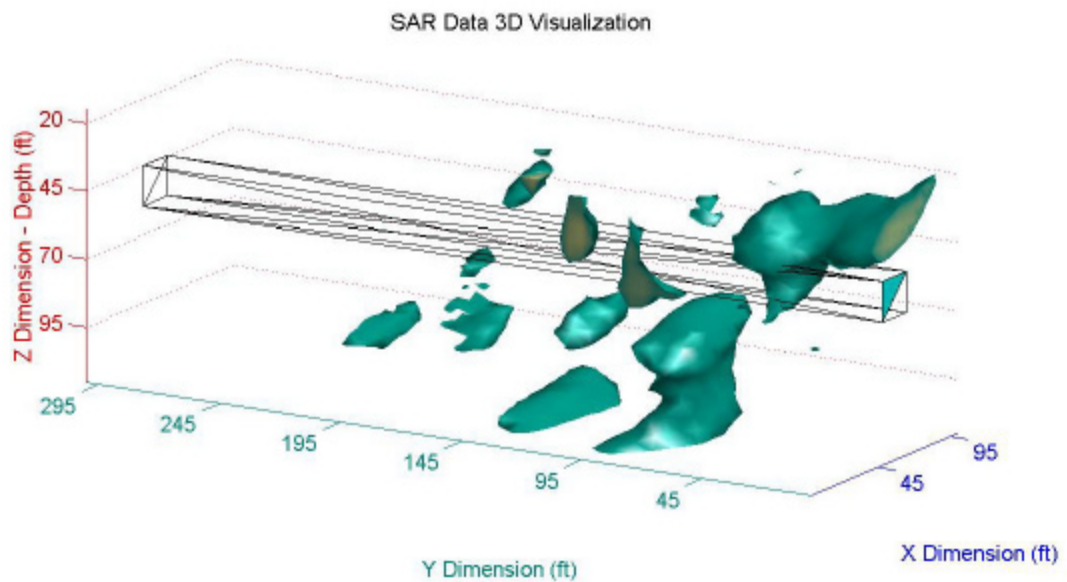
Figure 6.19 RMS Reduced NTS Scene, 100x300 Array

6.3.2.2 NTS with 100x100 Receiver Array

Performing the RMS-subtraction algorithm to the data set of the sub-array yields a similar result to that of the mean-subtraction technique. Figure 6.20 illustrates the effect. A very strong return from the tunnel location and a small patch of clutter is produced in the ensuing image.
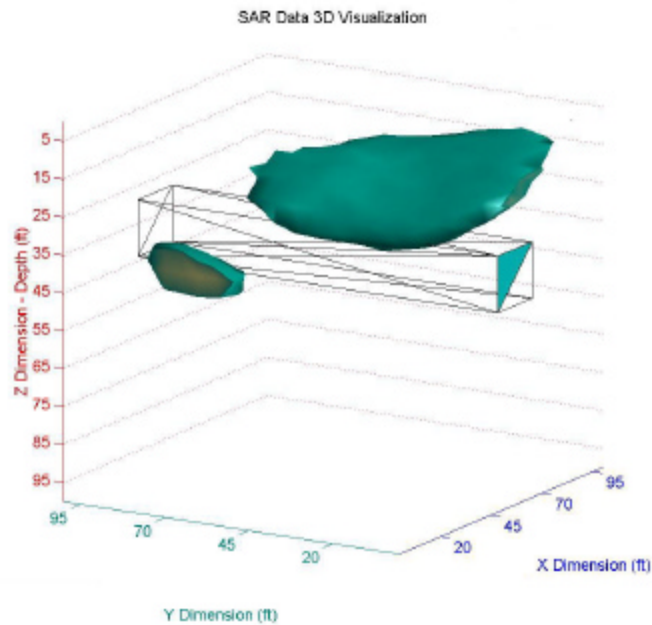
Figure 6.20 RMS Reduced NTS Scene, 100x100 Array

6.3.2.3 ZCA

The performance of the RMS-subtraction algorithm in the ZCA scene is very similar to that of the other two scenes. The result is also very similar to that of the mean-subtraction filter presented in section 6.3.1.3. Figure 6.21 gives an illustration of the outcome. Some of the clutter is eliminated, but the target is attenuated to the point of almost completely disappearing.

SAR Data 3D Visualization
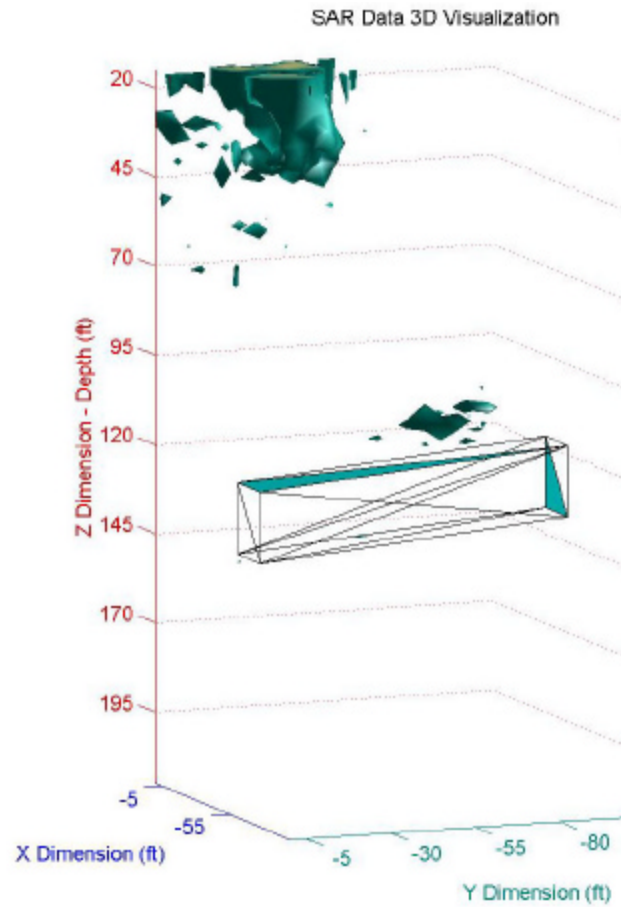
Figure 6.21 RMS Reduced ZCA Scene

6.3.2.4 RMS Window Filters

A three-dimensional window with equal sides was chosen for the RMS-subtraction algorithm. It extends across 3-voxels in each direction. The VUT is the voxel directly in the center of the window. Figure 6.22 shows the image resulting from applying this filter to the full-array of NTS data. The clutter is nearly completely eliminated, but so is the target.

Applying this filter to only the first sub-array of the data from the NTS scene results in the image of Figure 6.23. Although much of the signal is in the general area of

the tunnel location, it's attenuated to the point where discerning a target is difficult.
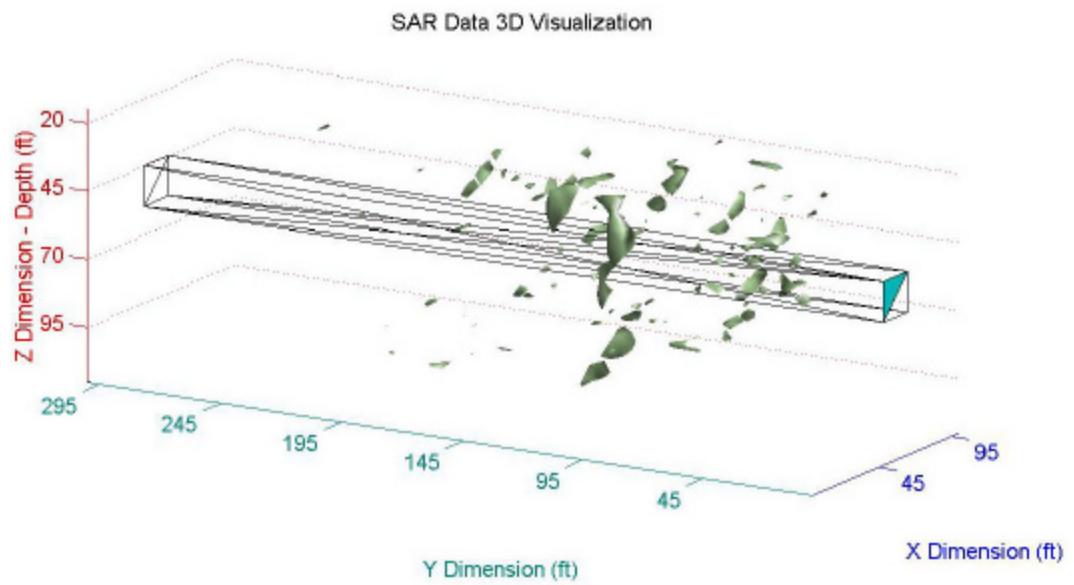
Figure 6.22 RMS Window Reduced NTS Scene, 100x300 Array
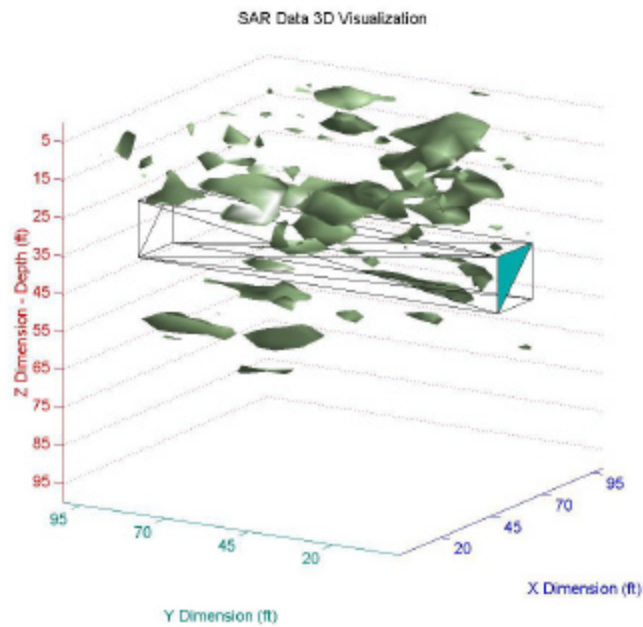
Figure 6.23 RMS Window Reduced NTS Scene, 100x100 Array

Figure 6.24 illustrates the result when this filter is applied to the data of the ZCA

scene. A good portion of clutter remains near the surface. The signal located in the

vicinity of the tunnel is attenuated such that detection of the target is difficult.
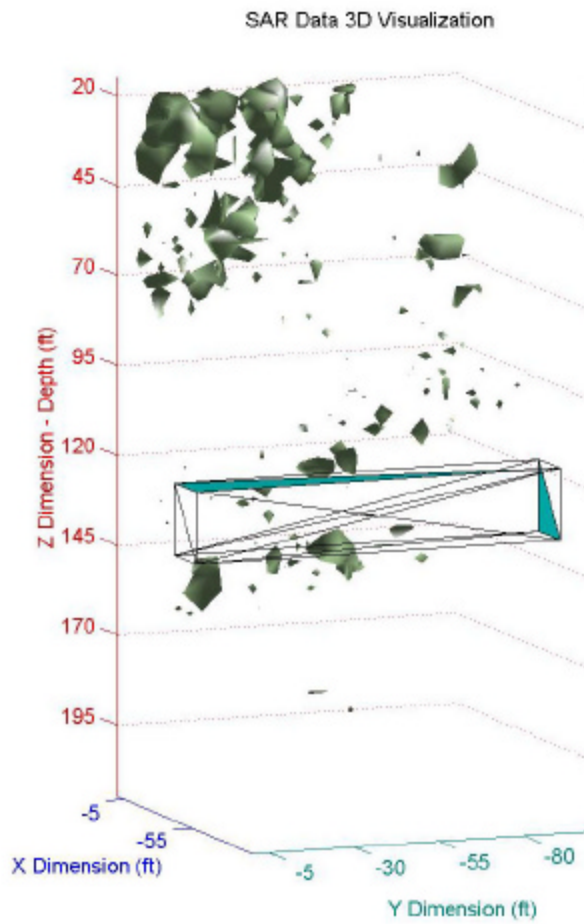
Figure 6.24 RMS Window Reduced ZCA Scene

The RMS method produced similar results to that of the mean method presented

in section 6.3.1. Table 6.2 summarizes the metric results for the RMS technique. Again,

the total RMS technique did not perform well in the full-array NTS or ZCA data set.

However, it improvements in the ratio of the target magnitude to clutter magnitude was

realized from the unfiltered cases.

The RMS window technique produced an improvement in the target magnitude to

Method	Full-Array of NTS Data	First Sub-Array of NTS Data	ZCA Data
Unfiltered Data	.2155	.2693	.1576
Total RMS	.2034	.3064	.1096
RMS Windows	.3100	.3844	.1676

Table 6.2 Metric Summary for RMS Techniques

clutter magnitude in all three data sets, even though some of the targets are not visible in some of the resulting images. The improvement noted in Table 6.2 is not on the order of the improvement seen in Table 6.1 of the mean window techniques. This can be attributed to the size of the window used in this filter.

This is an example illustrating the effect of not selecting an optimal window size for the sliding window. The targets in every scene are attenuated to such a degree that distinguishing them as targets in the images is nearly impossible, although the target magnitude to clutter magnitude ratio is improved from the unfiltered data set. The implications of window sizes will be discussed in the next chapter.

6.3.3 kNN

6.3.3.1 NTS with 100x300 Receiver Array

The Sobel edge detection algorithm of section 6.2.1 gave us three possible targets. The first possible target had traversed 10 voxels in the x-direction, 13 voxels in the y-direction, and 10 voxels in the z-direction. This results in a target that expends 1300 voxels. As mentioned in section 5.3.3, the k-number must be chosen large enough such that it doesn't cause the target to eliminate itself. To prevent the target from reducing itself, the k-number should be larger than the target. It was decided to use a k-number that is twice as large as the target. This is large enough to prevent the target from

eliminating itself, but not so large that the processing time becomes unreasonable. That

gives us a *k*-number of 2600 voxels in this case. The algorithm scrolls through each

voxel, identifying the 2600 voxels nearest in value to the current VUT. The mean of

those voxels is then subtracted from the VUT. Using a *kNN* filter with these parameters

on this data set results in the image displayed by Figure 6.25. The edge detector defined

two other possible targets existing in this scene. Those targets contain 720 and 1848

voxels, respectively. Figure 6.26 gives the image after applying this filter with a *k*-

number of 1440 and Figure 6.27 displays the result of this filter with a *k*-number of 3696.

The resulting images are very similar. A large return at the start of the tunnel location is

evident, but the remaining clutter makes detection of the target a challenge. It should not

be surprising that the filter with its *k*-number equal to 1440 (Figure 6.26) attenuates the

most clutter, because it uses the fewest *k*-nearest neighbors. It should also be noted that

this same filter results in the smallest target reflection from the tunnel location.
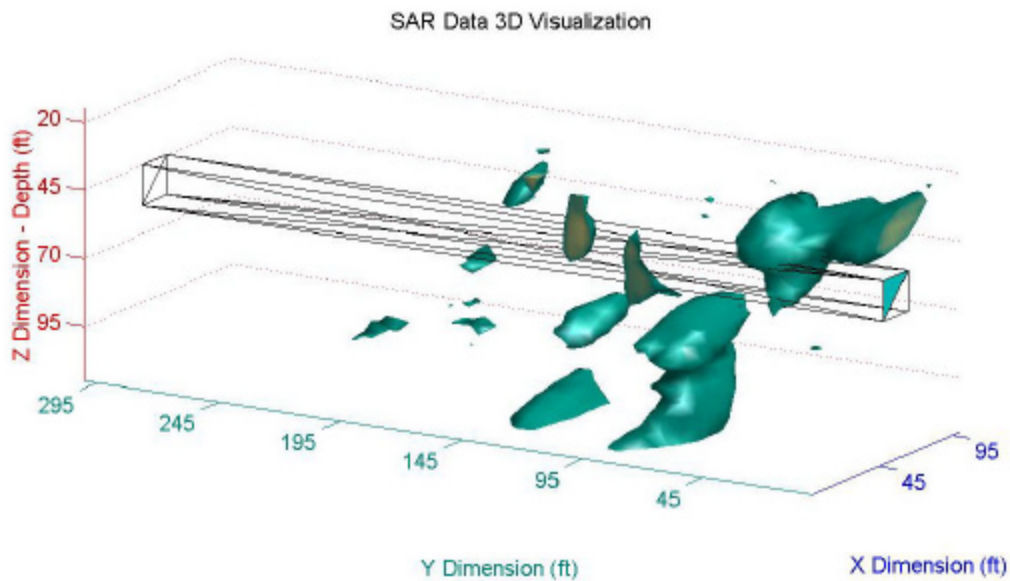
Figure 6.25 *kNN* Reduced NTS Scene, 100x300 Array, *k* = 2600
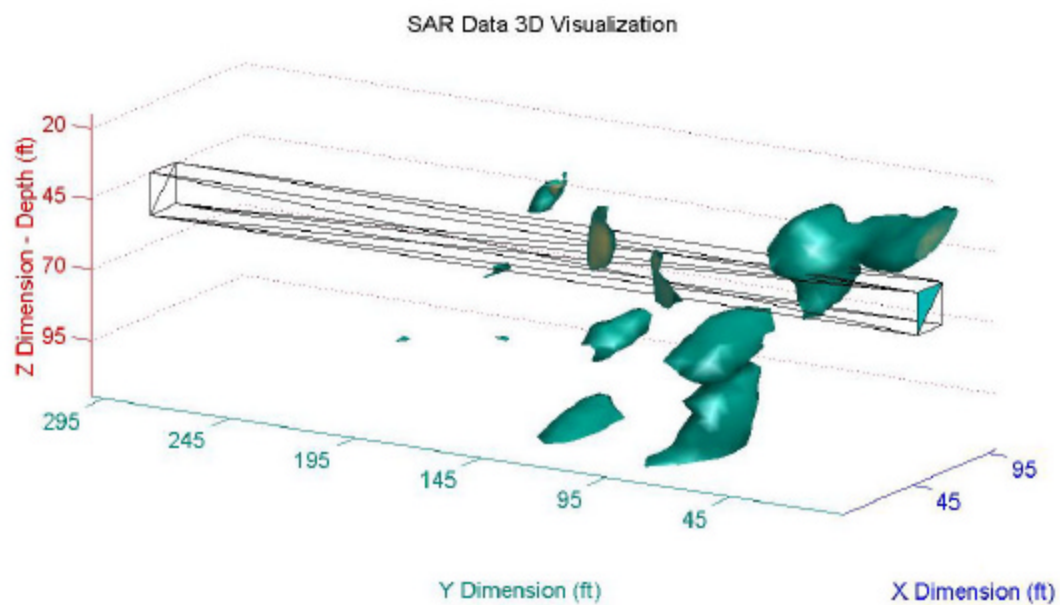
Figure 6.26 *kNN* Reduced NTS Scene, 100x300 Array, $k = 1440$
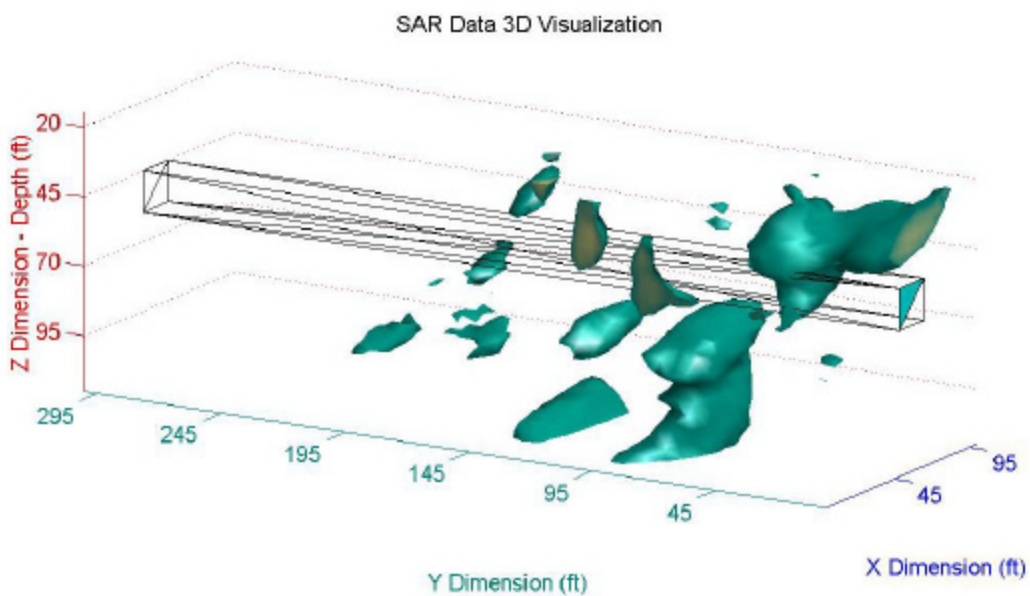
Figure 6.27 *kNN* Reduced NTS Scene, 100x300 Array, $k = 3696$

In section 6.2.1 the size of the suspected target in the NTS sub-array data set was found. The edges result in a target that measures 9 voxels in the x-direction, 15 voxels in the y-direction, and 5 voxels in the z-direction. That gives a target that traverses 675 total voxels. Calculating the k-number for this target generates a k-number of 1350. The resulting image is shown in Figure 6.28. A strong signal is detected in the in the area immediately above the actual location of the tunnel. A small residual clutter return remains in the reduced scene.
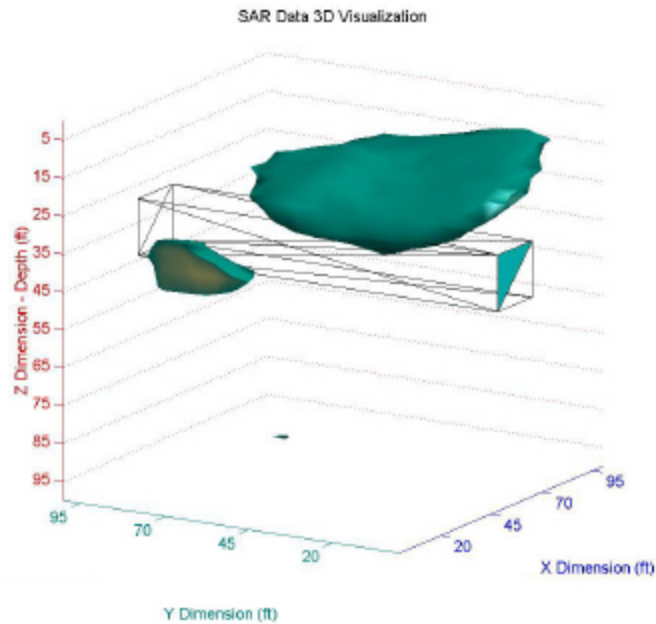
Figure 6.28 *kNN* Reduced NTS Scene, 100x100 Array, k =1350

Figure 6.29 illustrates the effect on the scene when the k-number is not selected correctly. It is an image from the same 100-feet by 100-feet sub-array scene of the NTS data, reduced with a k-number of 5. As can be seen, the target signal is much more attenuated than the example given in Figure 6.28 above and the only difference between

the two images is the selection of the k-number.

SAR Data 3D Visualization
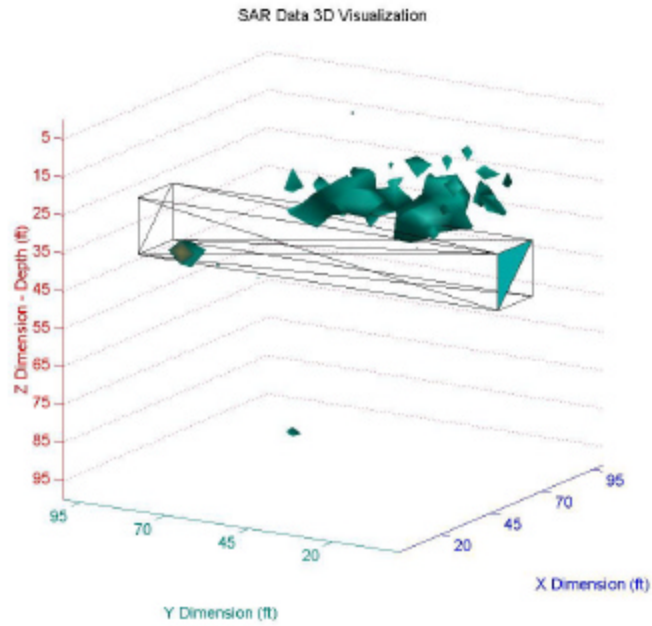
Figure 6.29 *kNN* Reduced NTS Scene, 100x100 Array, $k = 5$

6.3.3.3 ZCA

The edge detection algorithm found two possible targets in the ZCA scene. Section 6.2.1 defined the targets such that they include 616 and 968 voxels, respectively. The method of finding the k-number produces a value of 1232 for the first target and 1936 for the second target. The resulting images are shown in Figures 6.30 and 6.31. They are very similar.   Again, the effect the value of the k-number has on attenuation is seen across the scene. The smaller k-number results in more attenuation. Also in both images, the reflection originating from the real tunnel location experiences so much attenuation that it no longer resembles the tunnel.

SAR Data 3D Visualization
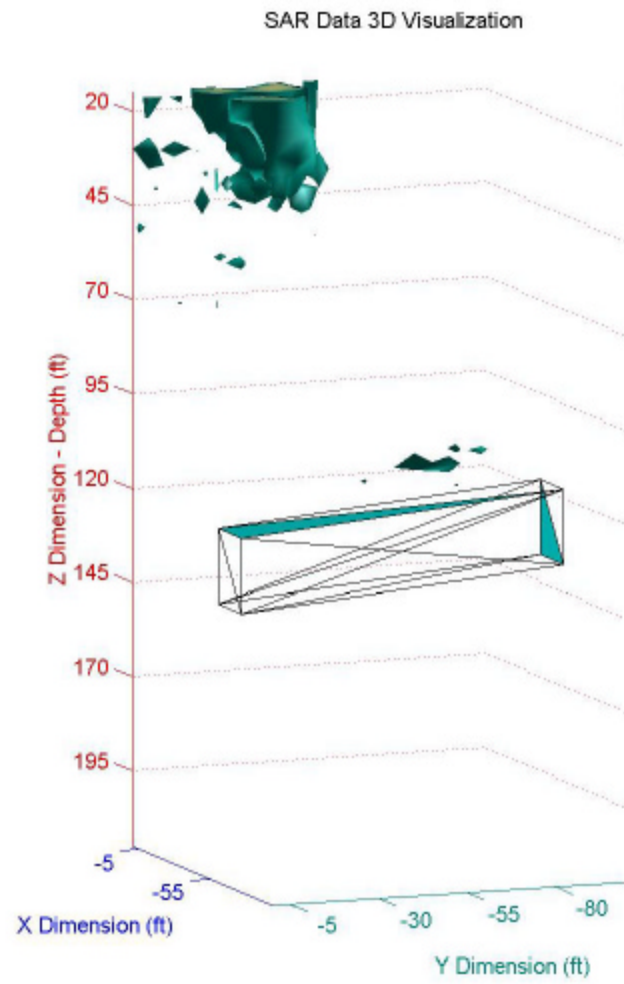
Figure 6.30 *kNN* Reduced ZCA Scene, *k* =1232
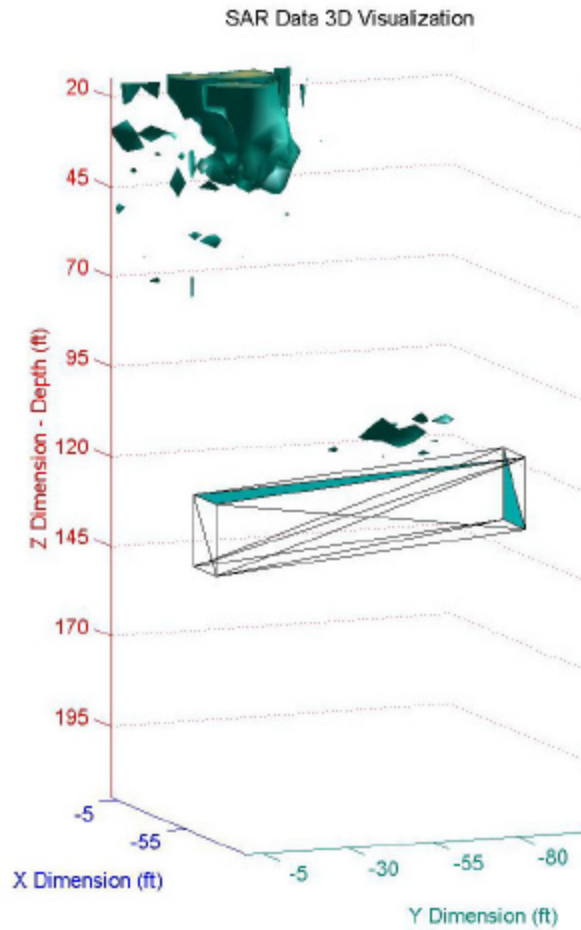
Figure 6.31 *kNN* Reduced ZCA Scene, $k =1936$

<u>*6.3.3.4 kNN Window Filters*</u>

The window selected for the *kNN* algorithm was a three-dimensional window extending 7-voxels in each dimension. The voxel in the center of the window is selected as the VUT. The *k*-number was defined as half of the voxels contained by the window. This was determined through some heuristic analysis and experimentation. The maximum possible *k*-number that could be chosen is all of the voxels except the VUT. Yet, this would produce precisely the same filter as the mean-subtraction algorithm described above in section 6.3.1. The minimum possible *k*-number (in order to modify

the scene in some manner) that could be chosen is a single voxel. This would result in an image that was severely attenuated, because the single voxel closest in value to the VUT would be subtracted from the VUT. In an effort to prevent duplication of the mean-subtraction algorithm and severe suppression of the scene, the k-number was chosen to be equal to half of the voxels in the window. This seemed to attenuate the clutter adequately without allowing the target to completely eliminate itself. The image shown in Figure 6.32 is the result of applying this filter to the full-array of NTS data. The clutter is attenuated little more than the reflection immediately above the tunnel location.
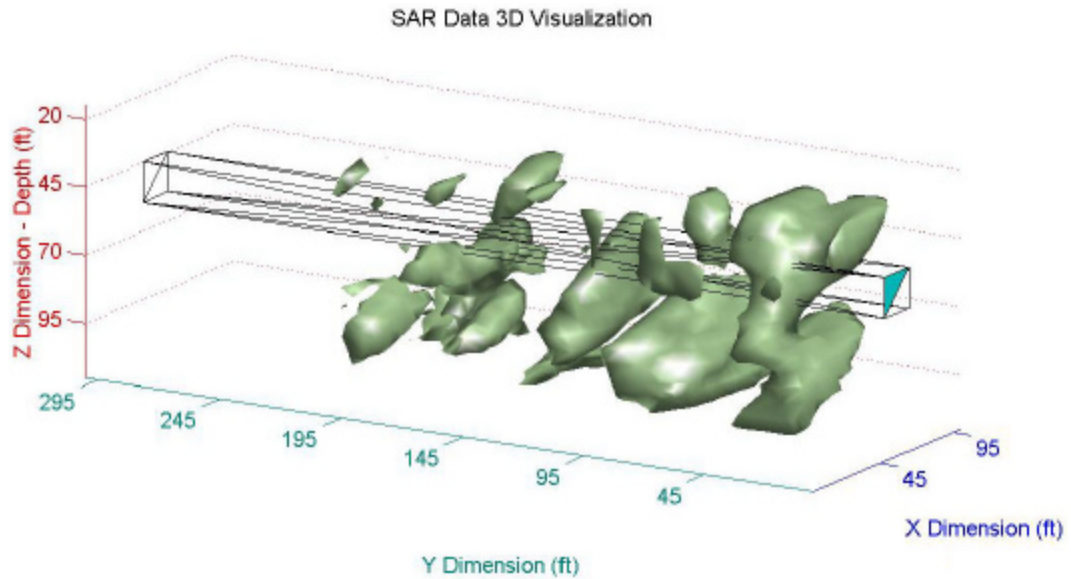
Figure 6.32 *kNN* Window Reduced NTS Scene, 100x300 Array

The result of employing this filter to the first sub-array of the NTS data is shown in Figure 6.33. While there is a large reflection originating from the tunnel location, enough clutter exists to cast doubt on the true location of the target.

Figure 6.34 illustrates the effect of applying this filter to the data of the ZCA

scene. The reflection from the tunnel location is very evident. The clutter near the
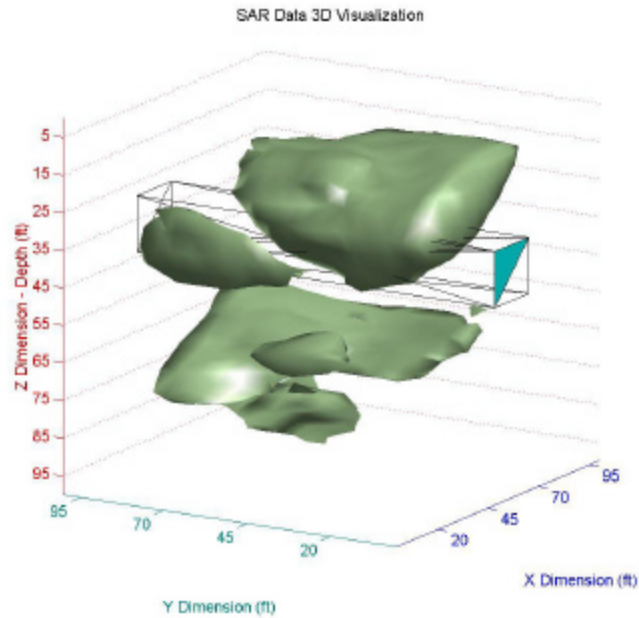
Figure 6.33 *kNN* Window Reduced NTS Scene, 100x100 Array

surface has been reduced significantly. This image certainly promotes the actual tunnel

location as the target. Additionally, notice how a very strong reflection at the location of

the second mine drift appears. This reflection is stronger in this example than in any

other scene, including the original.
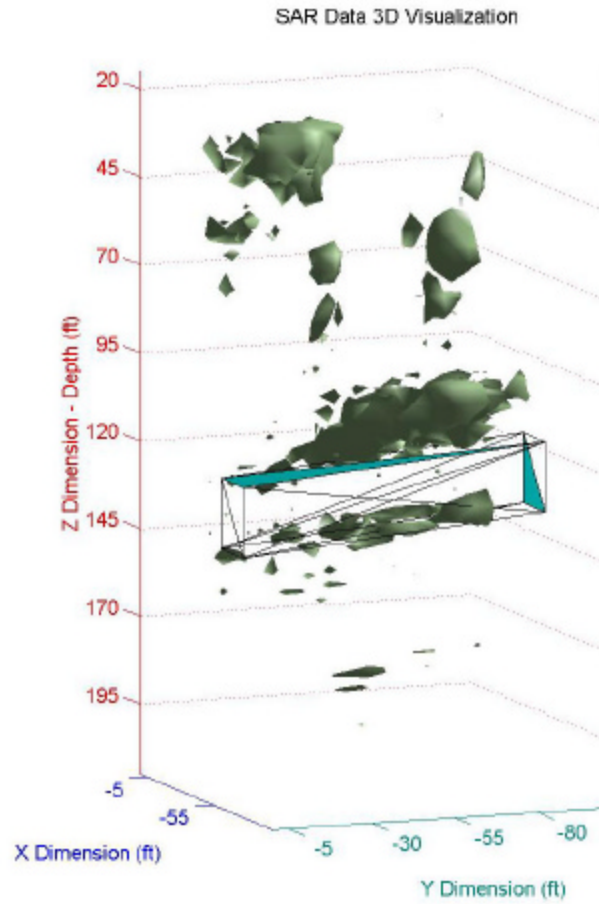
Figure 6.34 *kNN* Window Reduced ZCA Scene

Table 6.3 displays the metric results for the *k*-nearest neighbor algorithm. The

images are very similar to the other two estimation-subtraction techniques. Across

Method	Full-Array of NTS Data	First Sub-Array of NTS Data	ZCA Data
Unfiltered Data	.2155	.2693	.1576
kNN Total	.2473 (2600)	.4316 (1350)	.1328 (1232)
kNN Total	.2614 (1440)	.4190 (5)	.1872 (1936)
kNN Total	.2402 (3696)	N/A	N/A
kNN Windows	.3891	.5585	.2909

Table 6.3 Metric Summary for *kNN* Techniques

the same data set, the *k*-number doesn't appear to have a grand effect on the metric result. Even in the case where the *k*-number was 5 (Figure 6.29) the metric produced a similar target magnitude to clutter magnitude ratio. However, the images produced with small *k*-numbers were quite attenuated.

The 7-voxel cube window worked well to localize the statistics and produce good metric results and filtered images. In fact, the largest metric ratio was generated for the ZCA data set with this window.

6.3.4 Wiener Filter

6.3.4.1 Reducing Noise

The first experiment run with the Wiener filter is commensurate with the classical use of the filter. First, a scene that had one large target at a resolution of five feet was generated. The original scene is shown in Figure 6.35. Normally distributed, zero-mean random noise was then generated by the MatLab command *randn* and added to the original scene. The noise had a peak amplitude greater than that of the target. The resulting image is shown in Figure 6.36. This was used as the clutter term in the Wiener filter. To create a target reference, the target was moved from its location in the original scene to the center of the scene (Figure 6.37). The reference was put in the center of the scene as a default and to show that the Wiener filter is not affected by the spatial location of the reference. The reference does not need to match the true target location for the Wiener filter to successfully reduce the noise. Using Figure 6.36 as the clutter term and Figure 6.37 as the target reference term to generate the transfer function of the Wiener filter produces the restored output of Figure 6.38.
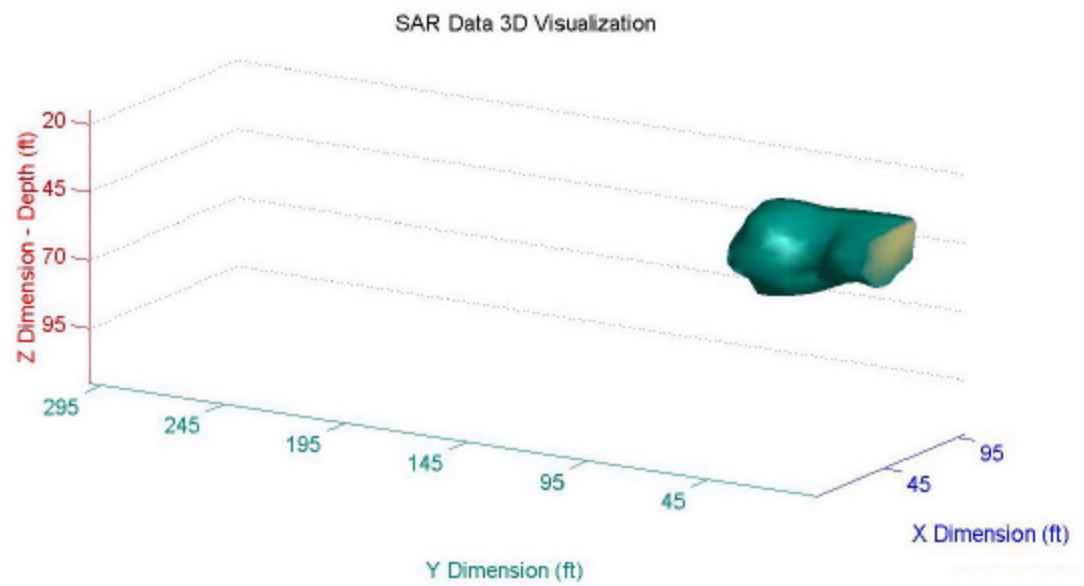
Figure 6.35 Original Scene for Classical Example
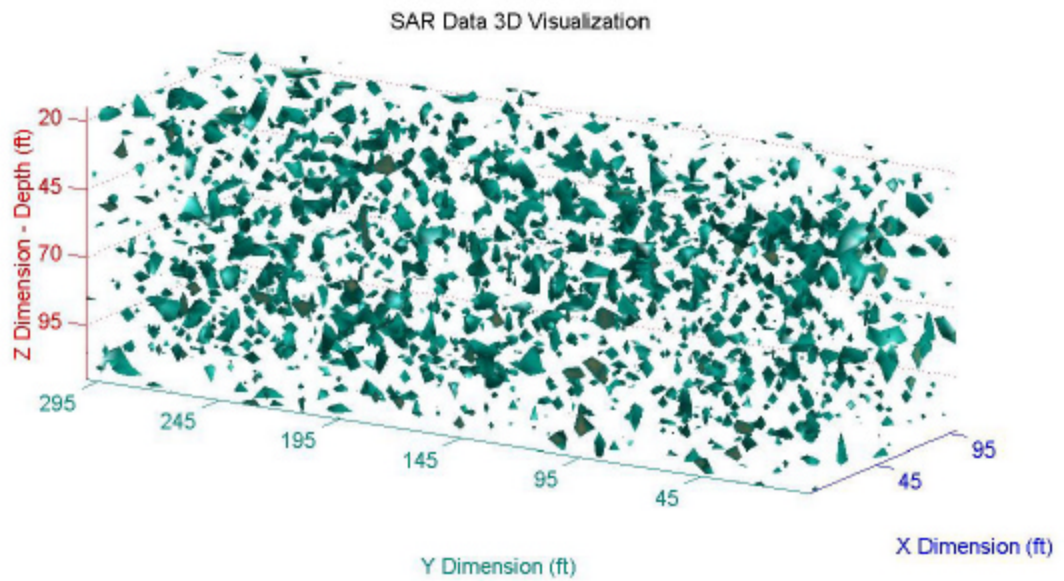
Figure 6.36 Original Scene with Random Noise Added
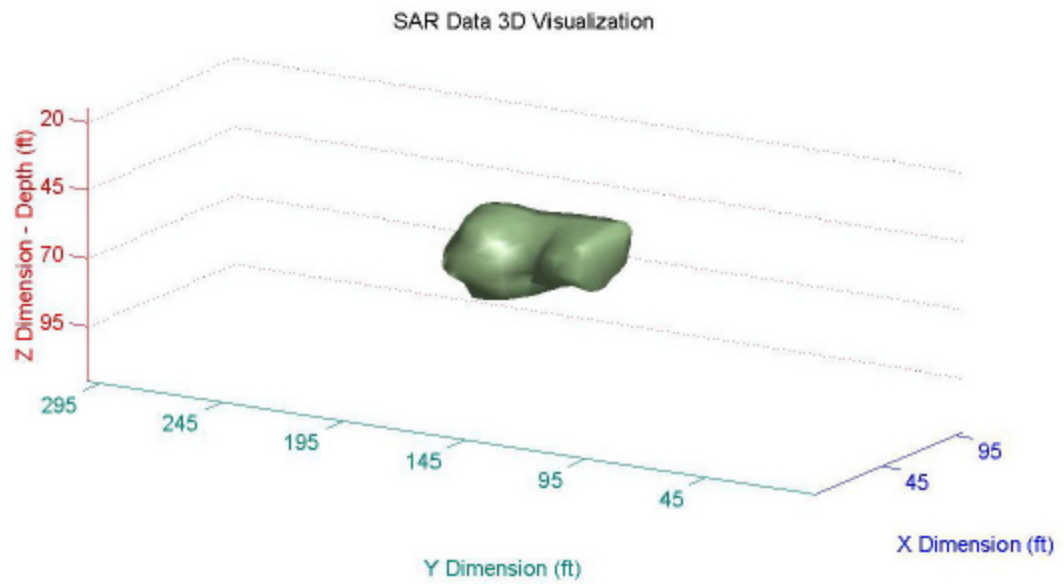
Figure 6.37 Target Reference for Classical Example
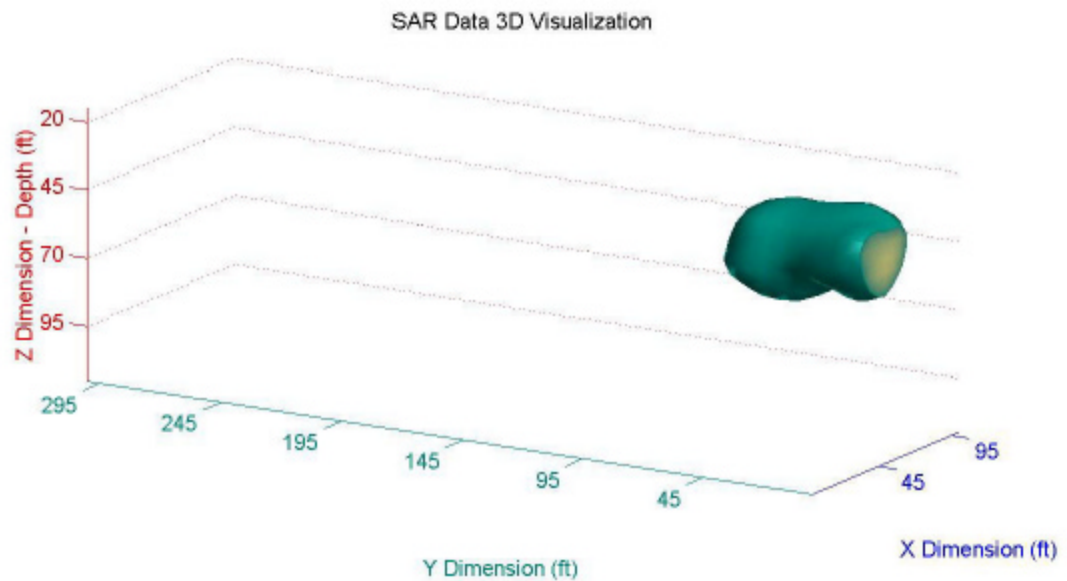
Figure 6.38 Wiener Filter Restored Scene

The Wiener filter has restored the noisy image to a near replica of the original and accurately identified the target location. This illustrates that the Wiener filter reduces

zero-mean noise very well, which has been shown by many other authors previously. Next, an investigation into how well it can reduce a real clutter environment that is spectrally similar to that of a real target will take place.

6.3.4.2 Reducing Real Clutter with a Ideal Target Reference

The performance of the Wiener filter in situations with real collected data will now be examined. Before using a simulated target reference, the target data contained in the original scene will be used. Since ground-truth of the tunnel location is available, a target reference will be formed with voxels in the immediate area of the real tunnel location. The target reference for the NTS full-array data needs to be defined first. If the scene is eliminated, except for the reflection immediately around the tunnel location, the target displayed in Figure 6.39 is left. The clutter term is defined with the data from the entire scene. This is given in Figure 6.40. Calculating the Wiener filter with this data set and reducing the data leaves the result in Figure 6.41. The tunnel appears larger than the reference, but is in the same location. The majority of the clutter is significantly
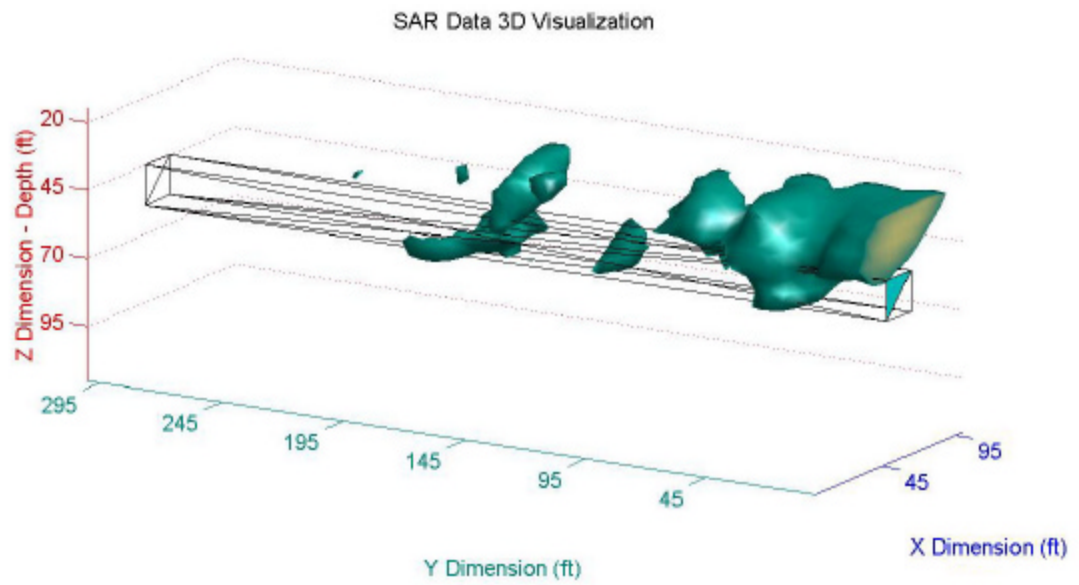
Figure 6.39 Ideal Target Reference NTS, 100x300 Array
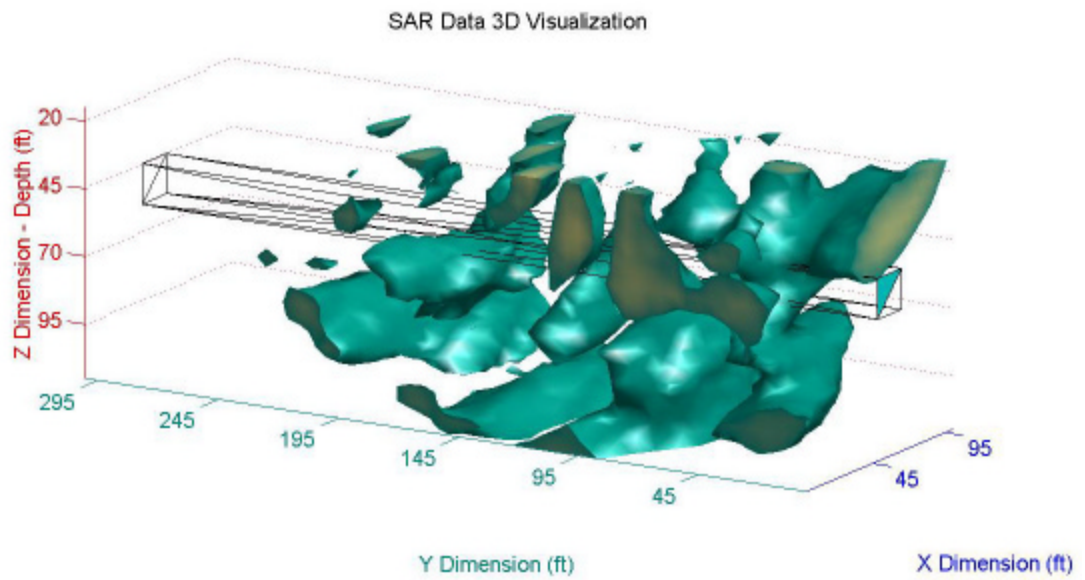
Figure 6.40 Clutter Reference – Entire Scene of NTS, 100x300 Array
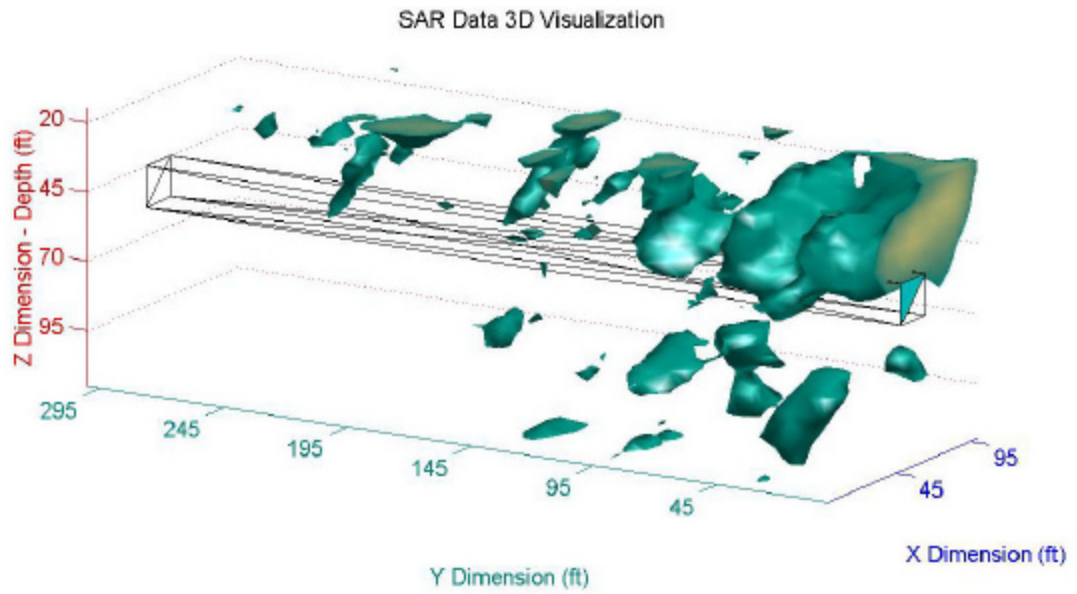
Figure 6.41 Wiener Filter Restored Scene

attenuated, leaving the tunnel return as the lone target of interest.

The ideal target reference for the first sub-array of the NTS data is given in Figure 6.42. Again, this was formed with the voxels in the immediate area of the actual tunnel location. The entire original data set is used as the clutter term. This scene is given in Figure 6.43. Using these terms and processing with the Wiener filter yields the reduced image of Figure 6.44.
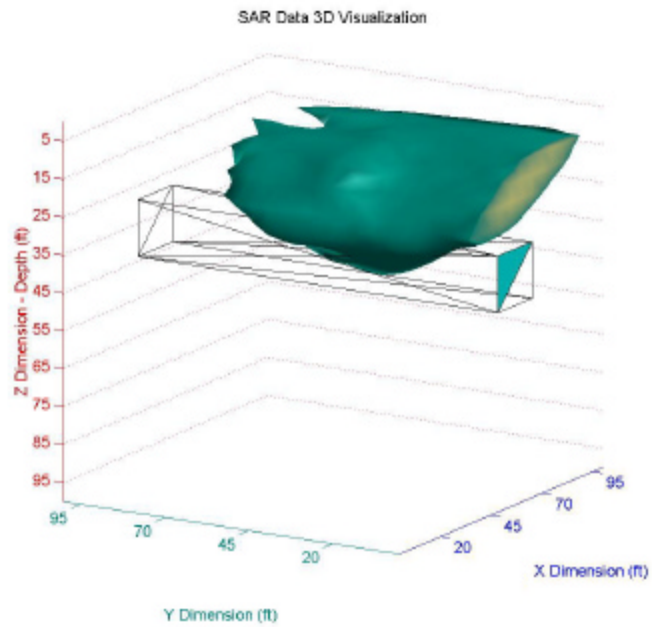
Figure 6.42 Ideal Target Reference NTS, 100x100 Array
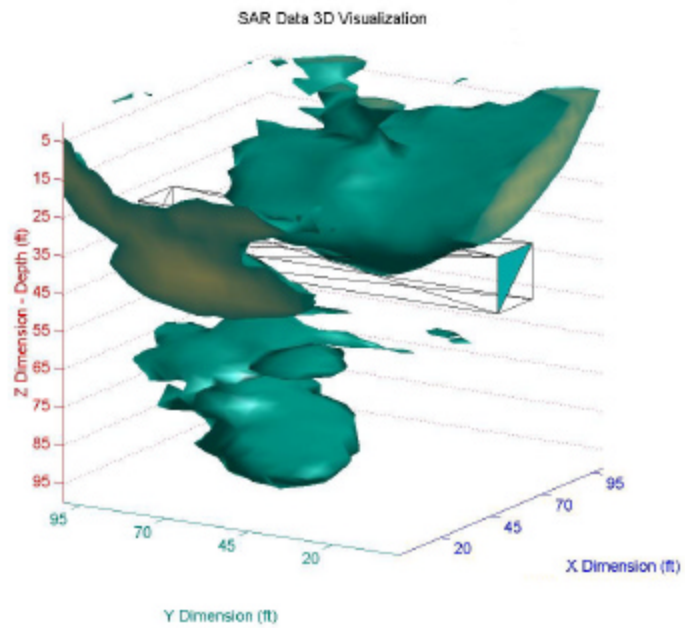
Figure 6.43 Clutter Reference – Entire Scene of NTS, 100x100 Array
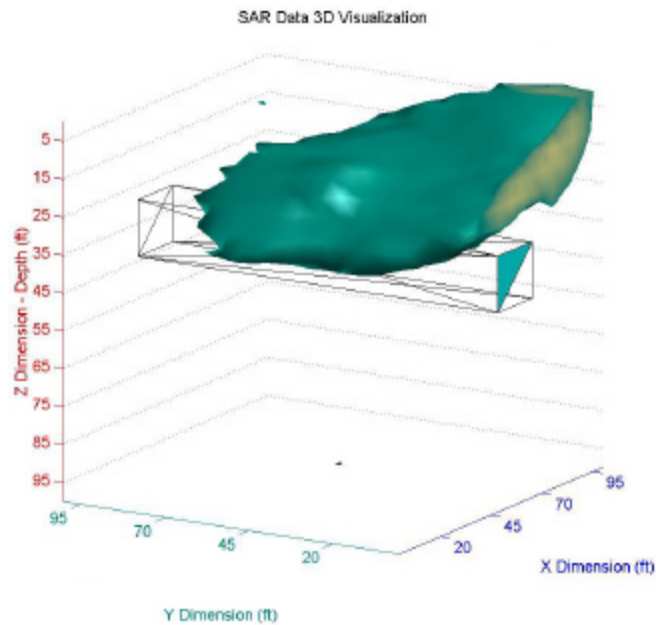
Figure 6.44 Wiener Filter Restored Scene

The target is slightly larger than the reference again, but is in the correct location. The clutter is attenuated to the point of being eliminated in the image.

The ideal target reference for the ZCA Hyatt mine is shown in Figure 6.45. This ideal target was formed with data voxels located in the immediate area of the actual tunnel. The clutter reference was created from the entire ZCA Hyatt data set and is given in Figure 6.46. These scenes create a Wiener filter transfer function that produce the scene in Figure 6.47.
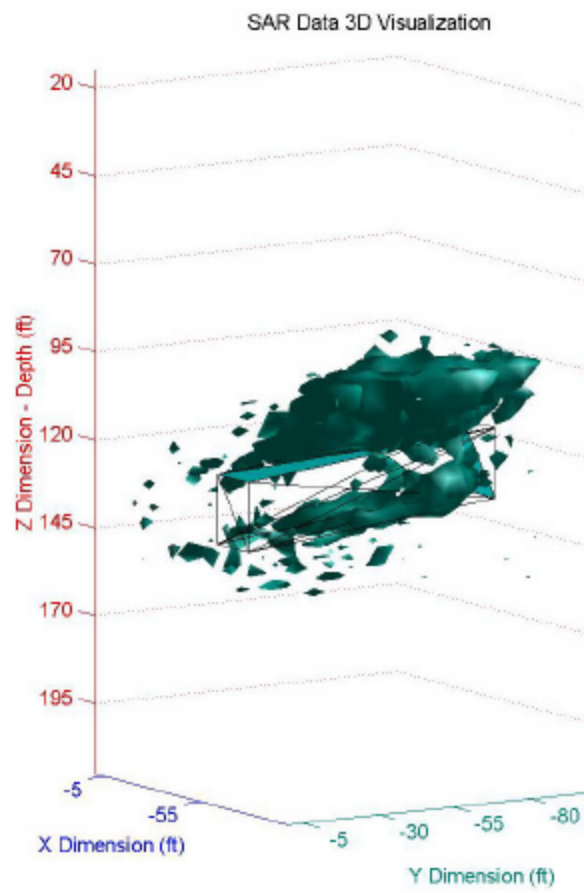
Figure 6.45 Ideal Target Reference ZCA
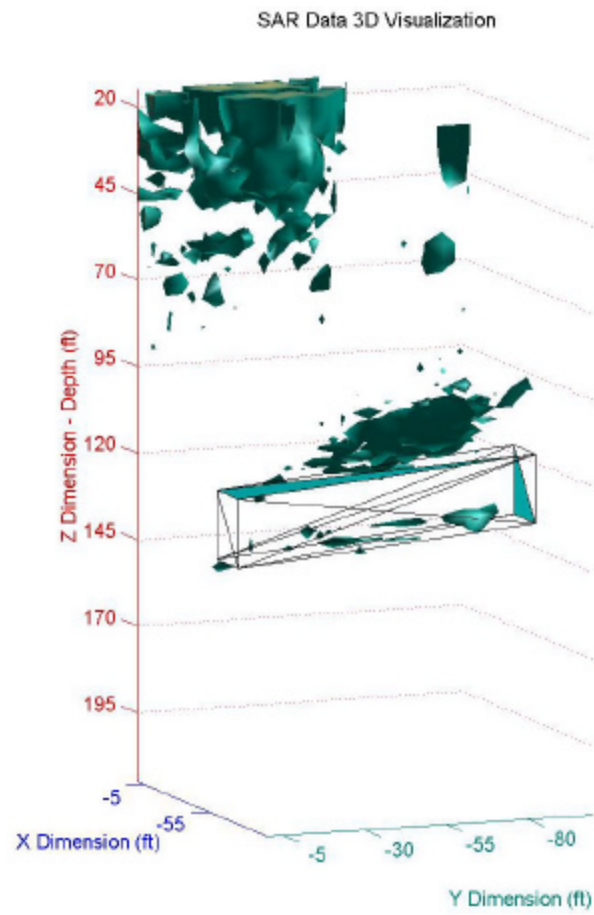
Figure 6.46 Clutter Reference – Entire Scene of ZCA
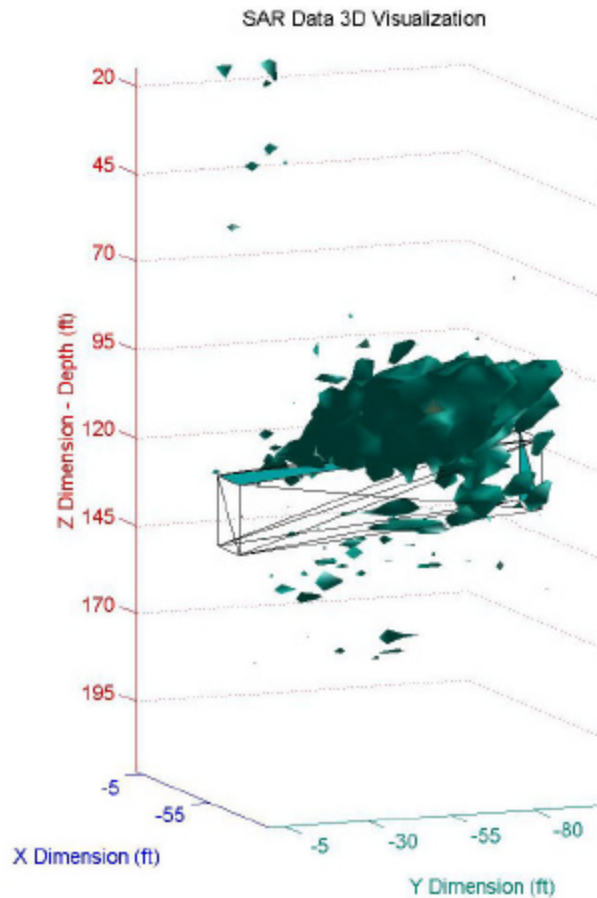
Figure 6.47 Wiener Filter Restored Scene

The target is larger than the reference, but exists in the correct location. A very small amount of clutter is visible in the resulting image.

This method clearly uses some *a priori* information regarding the target size and location. The target spectral characteristics were determined from the data existing in the target location of the actual scene. So, the results seen here are nearly the best that can be achieved when using the Wiener filter in a clutter environment. Optimizing the target data would create the ideal target reference and generate the best results. The target data was not optimized in the scene. There was, however, a constraint imposed not to require

a priori information in the schemes. So in the next section, restored scenes will be generated with information obtained from the original data and a simulated target model.

6.3.4.3 Reducing Real Clutter with a Simulated Target Reference

The Wiener filter algorithm will now be used with a simulated target reference based on the results of the edge detection algorithm. The clutter term of the filter will be generated with data from the original entire scene of the respective collect. In this section, the reference data will be created with the target simulator discussed in section 5.4. The target reference was normalized to the clutter data, so that they would share the same peak magnitude.

The Sobel edge detector found three potential targets for the full-array of the NTS data set. One will be chosen to illustrate the performance of the Wiener filter used in this capacity. The reference will be generated based on the target that stretches from 40 to 85 in the x-direction, 0 to 60 in the y-direction, and 10 to 55 in the z-direction. Figure 6.48
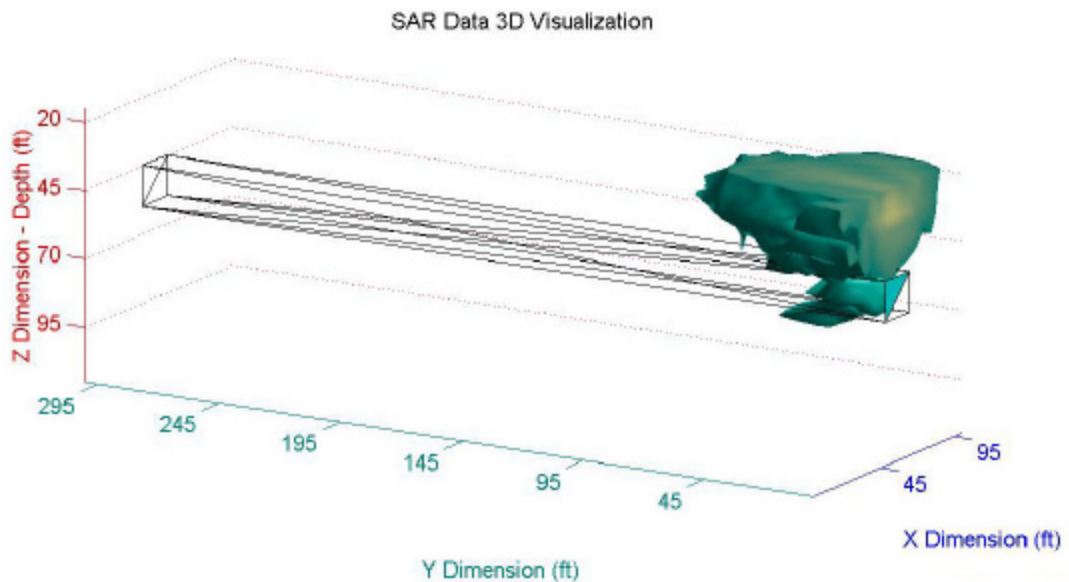
Figure 6.48 Simulated Reference for NTS target, 100x300 Array

113
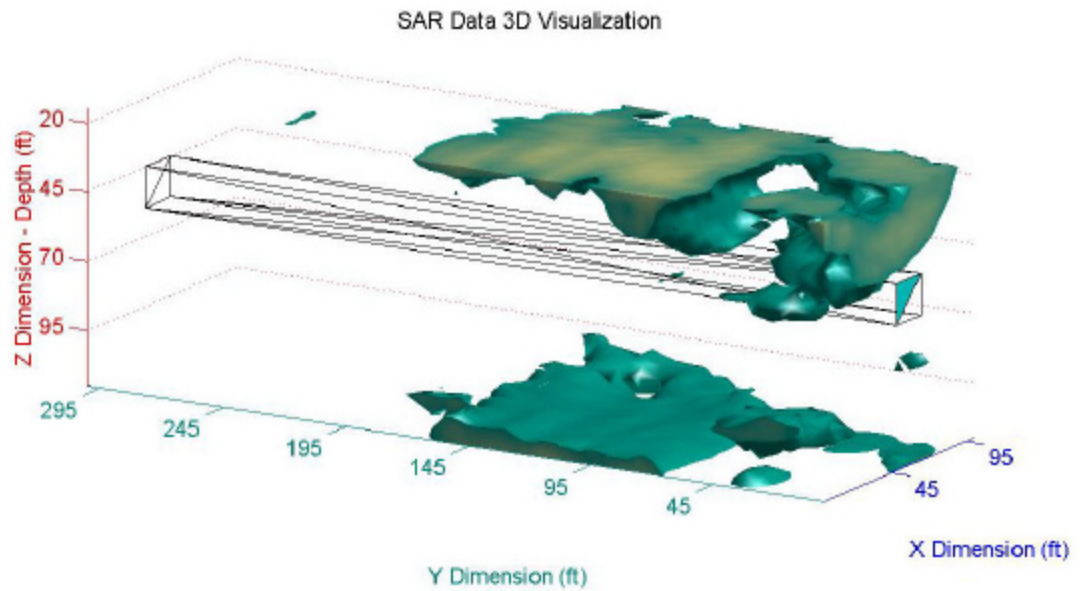
Figure 6.49 Wiener Filter Restored Scene, NTS 100x300 Array

illustrates the simulated target reference. The clutter reference in this trial is the entire

scene given in Figure 6.40. Figure 6.49 shows the Wiener filter restored image.

The resulting image has a large return emanating from the beginning of the tunnel

location, but large clutter returns above and below the tunnel are also existent.

The target defined by the Sobel edge detector for the first sub-array of the NTS

data had the following dimensions: 40 to 80 in the x-direction, 0 to 70 in the y-direction,

and 10 to 30 in the z-direction. The simulator produced the target reference that appears

in Figure 6.50. The clutter reference was generated from the data in Figure 6.43. The

restored image is given in Figure 6.51. It is much larger than the target or the real tunnel,

but is centered on the actual tunnel location. A small patch of clutter is evident on the

bottom of the scene.
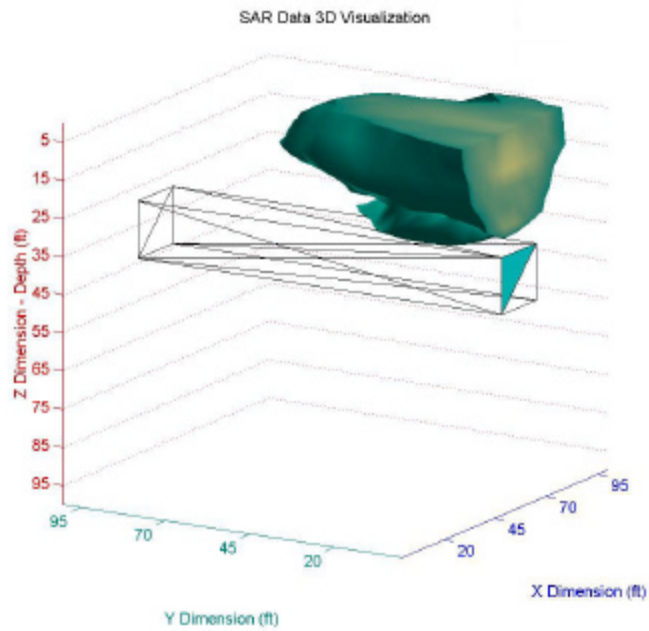
Figure 6.50 Simulated Reference for NTS target, 100x100 Array
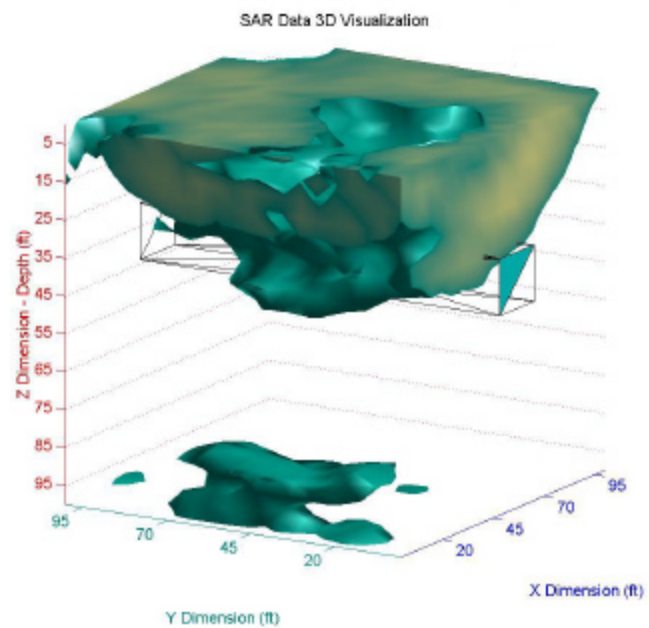
Figure 6.51 Wiener Filter Restored Scene, NTS 100x100 Array

115

The simulated target for the ZCA stretches from -40 to -80 in the x-direction, -30 to –80 in the y-direction and 95 to 125 in the z-direction. The image created by the simulator is shown in Figure 6.52. The clutter reference is created from the entire scene of data illustrated in Figure 6.46. The resulting filtered image is shown in Figure 6.53.
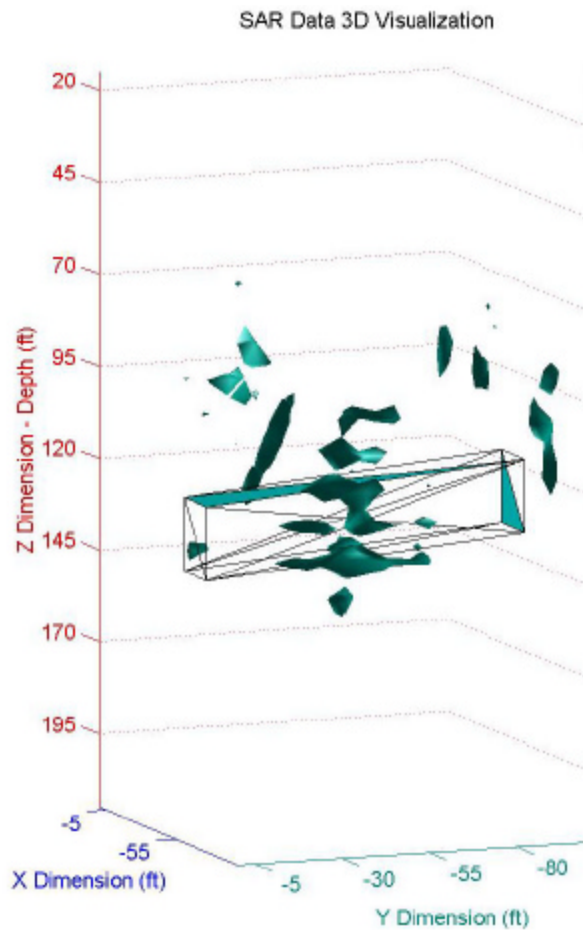
Figure 6.52 Simulated Reference for ZCA

The restored image is all clutter. Very little signal appears in the area of the expected tunnel return. This result can be traced to the target reference not modeling the real target correctly. This will be discussed in the following chapter.
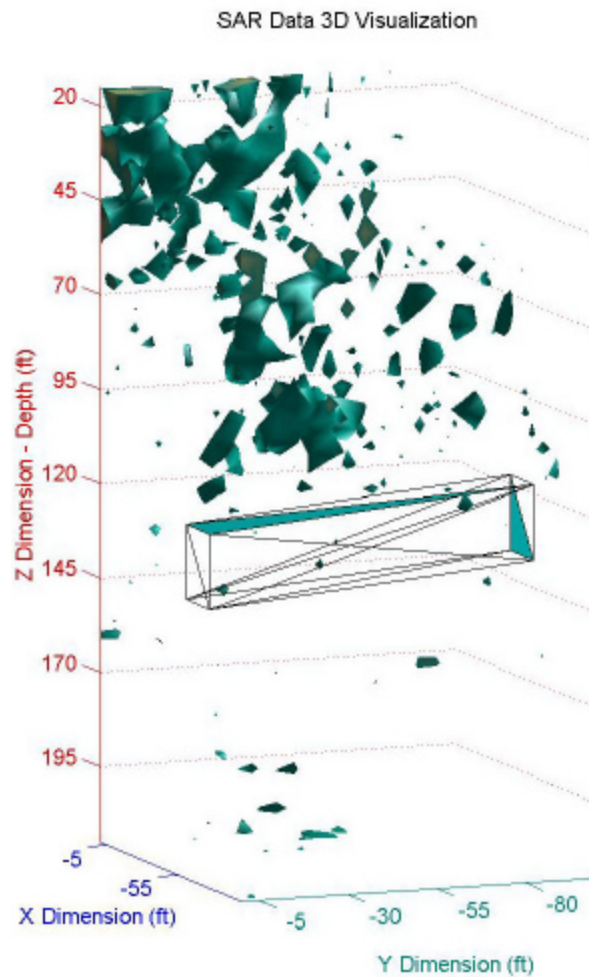
Figure 6.53 Wiener Filter Restored Scene, ZCA Scene

A summary of the metric data for the Wiener filter examples is presented in Table 6.4. The classic reduction of noise with the Wiener filter was shown in Figure 6.38. The ratio of target signal magnitude to clutter signal magnitude was improved by over 500%. This was by far the best result achieved by the Wiener filter. The Wiener filter was reliable in detecting the target when the target reference was generated with real target data. The full-array of NTS data does not have as impressive of a metric ratio as the other two data sets. It is believed that this is a product of the collection environment.

117

The receiver array was located on a hill, where the pulses received at the end of the array

had a two-way path length that was to 200 feet further to reach the target than the pulses

received at the beginning of the array. Including the receiver locations at the end of the

array would add very little target signal, for the amount of near-in clutter signal received.

The simulated target references had mixed results. A lot of the clutter signal was still

evident in all of the data sets after filtering. The ZCA data showed a slight improvement

in the target signal magnitude to clutter signal magnitude ratio and the full-array of NTS

data ratio showed no improvement. The sub-array of NTS data showed a moderate

improvement from the gains experienced with other techniques. Given the results of the

ideal target reference, it is believed the poor results of the simulated reference are a result

of inaccurately modeling the target.

Method	Classic Noise Example	Full-Array of NTS Data	First Sub-Array of NTS Data	ZCA Data
Unfiltered Noisy Data	.0396	N/A	N/A	N/A
Noisy Target	.2056	N/A	N/A	N/A
Unfiltered Ideal Data	N/A	.3083	.3971	.3860
Ideal Target	N/A	.4541	.6571	.7971
Unfiltered Simulated Data	N/A	.2155	.2693	.1576
Simulated Target	N/A	.2147	.3600	.1665

Table 6.4 Metric Summary for Wiener Filter Techniques

Chapter 7

Discussion and Conclusions

This research applied classical image processing techniques to wide bandwidth ground penetrating radar data in order to reduce the clutter of three-dimensional scenes. The clutter reduction techniques of this research were generally successful in increasing ratio of the target signal magnitude to the clutter signal magnitude. Tables 7.1 – 7.3 summarize the results of the techniques presented in chapter 6 by data set. The results

Method	Assumptions	Result	Reasons
Total Mean	Clutter with homogeneous statistics across entire scene	Target and clutter evident	Collection method added clutter to scene
Sliding Mean	Clutter statistics homogeneous in smaller local areas	Clutter reduced, target enhancement small	Window size eliminates clutter and some target signal
Total RMS	Clutter with homogeneous statistics across entire scene	Target and clutter evident	Collection method added clutter to scene
Sliding RMS	Clutter statistics homogeneous in smaller local areas	Less target magnitude than sliding mean, no definitive target in image	RMS subtracts more magnitude from image, poor window size selection
Total *kNN*	Spatially separated clutter with similar statistics	Target and clutter evident	Collection method added clutter to scene
Sliding *kNN*	Small pockets of spatially separated clutter with similar statistics	Clutter attenuation small	Window size too large, statistics vary
Wiener Filter	Dependant on modeling target reference statistics	Poor results when target reference created by simulator	Simulator does not model target spectral characteristics well

Table 7.1 Summary of Results, NTS 100x300 Array

of the research will be discussed in the remainder of this chapter. Also, areas of future

research for this work will be suggested.

Method	Assumptions	Result	Reasons
Total Mean	Clutter with homogeneous statistics across entire scene	Target detected in correct location, small clutter	Eliminated unnecessary adverse collection
Sliding Mean	Clutter statistics homogeneous in smaller local areas	Target detected, more target attenuation than total mean	Window restricts clutter involved in statistics
Total RMS	Clutter with homogeneous statistics across entire scene	Target detected in correct location, small clutter	Eliminated unnecessary adverse collection
Sliding RMS	Clutter statistics homogeneous in smaller local areas	Target attenuated, clutter reduced	Poor window size selection
Total *kNN*	Spatially separated clutter with similar statistics	Target detected in correct location, small clutter	Eliminated unnecessary adverse collection
Sliding *kNN*	Small pockets of spatially separated clutter with similar statistics	Target detected, large clutter	Window includes target into clutter statistics
Wiener Filter	Dependant on modeling target reference statistics	Best result for Wiener filter of all data sets, target grew larger, clutter reduced	Sobel edge detection defined target size too large

Table 7.2 Summary of Results, NTS 100x100 Array

Method	Assumptions	Result	Reasons
Total Mean	Clutter with homogeneous statistics across entire scene	Target severely attenuated, clutter slightly reduced	Strong clutter at surface, weak deep target
Sliding Mean	Clutter statistics homogeneous in smaller local areas	Target evident, clutter severely reduced	Window restricts statistics to local area
Total RMS	Clutter with homogeneous statistics across entire scene	Target severely attenuated, clutter slightly reduced	Strong clutter at surface, weak deep target
Sliding RMS	Clutter statistics homogeneous in smaller local areas	Severe clutter and target attenuation	Poor window size selection
Total *kNN*	Spatially separated clutter with similar statistics	Target severely attenuated, clutter slightly reduced	Strong clutter at surface, weak deep target
Sliding *kNN*	Small pockets of spatially separated clutter with similar statistics	Target enhanced, severe clutter attenuation	Window restricts statistics to local area
Wiener Filter	Dependant on modeling target reference statistics	Good results when actual target data use, poor results with simulated target	Sobel edge detection defined target size poorly, simulator poorly models target reference

Table 7.3 Summary of Results, ZCA

7.1 Image Enhancement Techniques

Sobel edge detection limits have an effect on how large the targets appear. The

edge detector has a profound effect on the performance of the techniques that use its

output to define the physical target characteristics. In this implementation, the target

parameters (size and depth) are determined by viewing the edge map. No analysis was

performed to determine what the optimal threshold should be to determine the target

parameters. Future work in this area should focus on developing an automated scheme to

define the parameters of possible targets in a consistent manner. Experimentation should

be performed to determine the optimal thresholds for a given target. An analysis with a constant false-alarm rate (CFAR) processor may help to determine this. The CFAR uses the statistics of the clutter around a given voxel to determine the threshold that should be used for detecting a target in a certain space. The CFAR would determine the varying threshold for a given false-alarm rate across an entire scene, which could lead to finding the optimal threshold for detecting a given target. This application could be implemented using the existing FIR local window processing.

The results of histogram equalization in this study were inconclusive. In part, that was due to the limited number data sets at the disposal of this research. In the data sets used, histogram equalization did not improve the selection of targets by the edge detector. Studies with controlled data sets where the ground truth is know would help determine the effectiveness of histogram equalization in GPR data.

7.2 Estimation-Subtraction

The first observation when analyzing the data was that the estimation-subtraction algorithms performed similarly when compared to one another. This is especially true comparing the mean and the RMS techniques. The improvement in the ratio of target signal magnitude to clutter signal magnitude was similar, relative to the unfiltered ratio. This is in part, because of the similar processing inherent in the estimation-subtraction filter that occurs and that there were no outlier data sets.

The second observation is that more receiver locations do not necessarily produce improved results. In nearly every case, the images produced by the first sub-array of NTS data gave better results than the images produced by the full-array of NTS data.

Also, none of the algorithms performed really well when operating on the full-array NTS data set. The primary reason is that the receiver locations at the end of the array contributed mostly clutter signal and very little target signal. In this case, adding more receiver locations added the amount of clutter in the scene. The collection environment attenuated the target after the first sub-array of Figure 6.2. Recall that the tunnel was only 30 feet deep at the beginning of the array, but 130 feet deep at the end of the array and that there was surface wetness. Although this system has shown the ability to penetrate to depths exceeding 130-feet, these conditions were not conducive to that task. The energy that reflected off of the target and was received by the last two sub-arrays would be much more attenuated than the energy received from the target in the first sub-array. This was noticed in analyzing the data, but correcting this problem requires *a priori* knowledge of the collection environment. Consequently, smart collection processes can help eliminate clutter and reduce the amount of clutter that needs to be suppressed by processing.

The estimation-subtraction algorithms performed differently in dissimilar environments. The techniques that estimated the clutter over the entire scene performed better in the NTS scenes; while the sliding window performed better in the ZCA mine scenes. This observation is primarily due to the depth of the tunnels. The NTS tunnel is relatively shallow when compared to the ZCA mine drift. The electromagnetic energy that reflected off of the ZCA mine drift was attenuated much more than the energy reflecting off of the large surface clutter in that scene. By virtue of propagating through less of the attenuating medium, the clutter magnitude was larger than the target

magnitude. The clutter was also represented by many more voxels than the target. These two factors contributed to the reduction of target signal in the ZCA images created by the techniques that estimated the clutter over the entire scene. The sliding local windows had better performance across the ZCA mine, because the local windows only consider those voxels in the window to calculate the clutter statistics. Therefore, the local windows use voxels in the immediate vicinity that experience similar propagation attenuation. The methods that estimated the clutter based on the entire scene performed well when the target magnitude was greater than or on the order of the clutter magnitude. The shallower tunnel of the NTS had this stipulation.

The size of the local windows has a profound effect on the resulting image. Windows that applied local statistics prevented large clutter sources from dominating the target in the filter. This was seen in the total estimation-subtraction methods. A 3-voxel by 3-voxel by 3-voxel window was used in the RMS sliding window case (section 6.3.2.4). This window attenuated the target and clutter signal across the entire scene in all three data sets. The estimation-subtraction filter depends on sampling across enough of the scene to mix some clutter in with the target signal. Otherwise, the local scene eliminates itself. A small window rarely provides adequate sampling. The windows need to be large enough to include some of the clutter in the statistics when over the target. The 7-voxel square and cube windows gave the best results. This reduces the attenuation of the target signal and prevents the target from eliminating itself. In the experiments presented here, the 7-voxel sided windows worked the best for deep targets. This is dependent on the size of the target and the size of the entire scene, however.

In the *kNN* estimation-subtraction technique, the *k*-number also has an effect on the resulting image. It was found that in deep targets, increasing the *k*-number resulted in an increased ratio of target signal magnitude to clutter signal magnitude. However, it doesn't have to be that large in order to be effective. In the NTS data set (where the target is shallow), however, the larger *k*-number resulted in a decreased ratio of target signal magnitude to clutter signal magnitude. When the *k*-number is small, the clutter is more susceptible to reduce itself. However, so is the target. When the *k*-number is large the reduction will generally be less in the target and clutter, because more voxels with statistics different than the VUT will be included in the data. There is tradeoff in selecting the *k*-number between reducing the clutter and reducing the target. The more alike the clutter is and the more different it is from the target, the better this algorithm will perform. The *k*-number was chosen such that it was twice that of the size of the target found by the edge detector and this was found to be a good value to use. This *k*-number was large enough to prevent the target from eliminating itself, but small enough to keep the processing time reasonable and keep the balance between the strong and weak targets. A small *k*-number will cause the target to reduce itself, as illustrated by Figure 6.28. This is caused when the *k*-nearest neighbors are similar to the value of the VUT.

7.3 Wiener Filter

Tan et al [13] state that the performance of the Wiener filter strongly depends on ability to match the exact spectrum of the clutter. While this is true, it would be prudent to add that the results presented here suggest the target reference is equally important. In this implementation, the clutter reference was defined as the entire scene of the collected

data. The results were a factor of how well the filter was able to match the target

reference. The Wiener filter performance drops rapidly when the target reference is not

the same as that used to construct the Wiener filter. This is illustrated by the difference in

restored images of sections 6.3.4.2 and 6.3.4.3. The more the reference matches the real

target signal, the better the filter performs. A similar problem would be experienced with

the clutter reference if the entire scene weren't used for it. In fact, Tan et al [13]

experienced this problem. The clutter reference used in this research matches the real

clutter very well, because the entire scene is used to calculate the clutter reference. So in

this method, duplicating the spectral characteristics of the target reference is the biggest

obstacle.

Generation of an accurate target model is crucial to the performance of the

Wiener filter. Future work should strive to improve the simulated target model. Work is

this area has already begun in the form of a GPR ray tracing model developed by Jeter

[24]. This simulator models the returns from extended targets by measuring the specular

reflection from the transmitter to the receiver location. It is suggested that future

simulated target work on the GPR system should form a hybrid of the completely diffuse

model used in this thesis and the specular model developed by Jeter. By merging the two

methods, a more realistic representation of the reflection caused by surface of the target

can be realized.

In an effort to illustrate the performance that can be attained when the target

reference is sufficient, an illustration of such a case with the model used in this research

will be shown here. A target reference that defines the ZCA target data more accurately

will be used. This target reference is given in Figure 7.1. The clutter reference is defined

as it was before, Figure 6.46. The result of the Wiener filter is given in Figure 7.2. The

clutter near the surface has been moderately attenuated and the target is much more

visible than the result found with the target reference defined by the Sobel edge detector.

The target signal magnitude to clutter signal magnitude ratio is .1655.
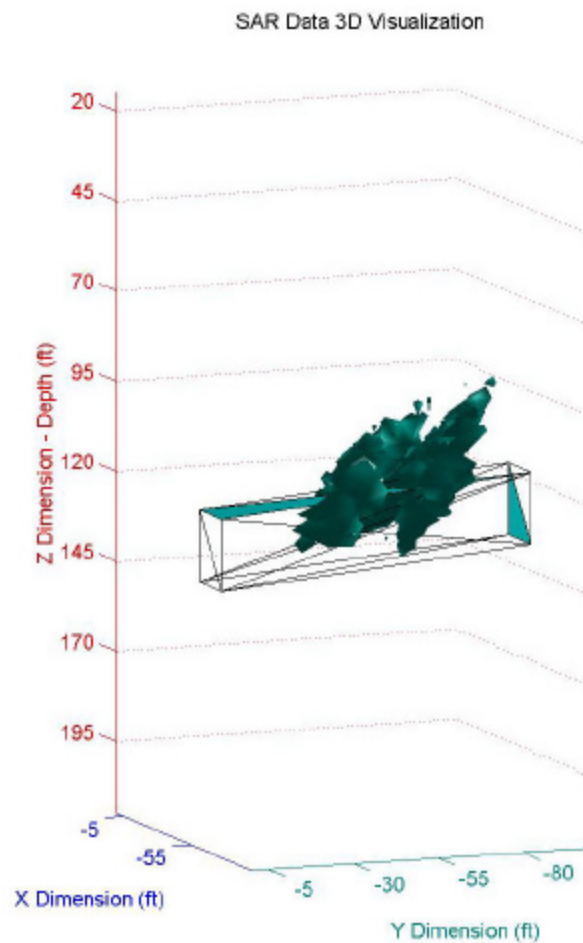
Figure 7.1 Simulated Ideal Target Reference, ZCA

126
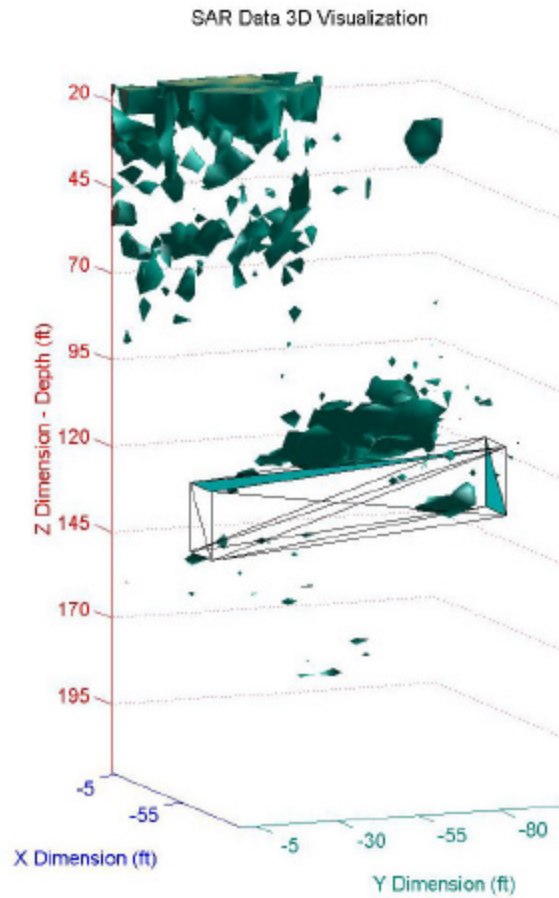
Figure 7.2 Wiener Filter Restored Scene, ZCA Simulated Ideal Target Reference

7.4 General Observations

The dielectric permittivity of the ground is a crucial variable in GPR. It can affect the apparent depth the target appears, the propagation of the electromagnetic energy, and the strength of the reflection. The data included in this thesis was collected without accurate measurements of the dielectric permittivity of the environment. So much of this work depends on this variable. Accurate ground dielectric permittivity models should be implemented. Controlled laboratory experimentation can help solve this problem, if the environment and target characteristics are known. In order for an accurate model of

clutter to be developed, extensive experimentation on the effects of dielectric permittivity in these radar frequencies must be conducted. Highly accurate models are necessary, because the spectral characteristics of clutter are so similar to that of the target. While subsurface terrain is usually given a range of 10-25 for the value of dielectric permittivity, a constant value of 17 was used in order to limit the variables in the experiments. While the results show this appears to be a good rough estimate, experiments must take place to verify the permittivity of each specific environment.

The reflectivity of the targets also needs to be researched. At this time it is uncertain as to what an echo signal from these targets should look like, given the parameters of this system. Nor is it fully understood what signal characteristics to expect. No controlled experimentation has taken place with known targets and known clutter environments. Experimentation should take place to determine the signal and visual characteristics of reflections from these similar structures. Ultimately the more that is known about the target and clutter environment, the better the filter parameters and simulation models used in detecting subsurface targets.

7.5 Summary and Conclusions

The effects that these image processing techniques have when applied to reducing the clutter of this GPR system has been illustrated. Estimation-subtraction filters have been used and the Wiener filter has been adapted to solve this particular problem. There was also a restriction implemented regarding the availability of *a priori* information to the filters. Image enhancement techniques have been used to determine the physical characteristics of potential targets in the scene.

Ultimately image processing techniques are judged on how well they can detect the target of interest. Although visual perception is usually used to determine if the target is in the scene, this is not a requirement. Perhaps the eye is not even the best way to determine if a target is present. The target signal magnitude to clutter signal magnitude ratio gave information about the existence of a target in some scenes where a target was not visible in the image. This information could be exploited to detect a target and identify where it exists.

Bibliography and Citation

[1]   John W. Brooks, Luc van Kempen, and Hichem Sahli, "Ground Penetrating Radar Data Processing: Clutter Characterization and Removal," IRIS-TR-0059, January 1999.

[2]   John Wayne Brooks, *The Detection of Buried Non-Metallic Anti-Personnel Land Mines*, Ph.D. Thesis, The University of Alabama in Huntsville, 2000.

[3]   L. van Kempen, H. Sahli, J. Brooks, and J. Cornelis, "New Results on Clutter Reduction and Parameter Estimation for Landmine Detection using GPR," Eighth International Conference on Ground Penetrating Radar, Gold Coast, Australia, May 2000.

[4]   John R. Schott, *Remote Sensing: The Image Chain Approach*, Oxford University Press, Inc, 1997.

[5]   M. Lundberg, "Reduction of Clutter in Infrared Images with Visual-Wavelength Measurements," IEEE 2000 International Geoscience and Remote Sensing Symposium, pp. 2377-2379, July 2000.

[6]   H. Brunzell, "Clutter Reduction and Object Detection in Surface Penetrating Radar," Proceedings of IEE Radar, Vol. X, no 449, pp. 688-691, October 1997.

[7]   Lennart Ljung, *System Identification: Theory for the User*, Prentice-Hall, Inc, Englewood Cliffs, New Jersey, 1987.

[8]   Merrill I. Skolnik, *Introduction to Radar Systems*, Second Edition, McGraw-Hill, Inc, New York, 1980.

[9]   Gopalan Ravichandran and David Casasent, "Minimum Noise and Correlation Energy Optical Correlation Filter," Applied Optics, Vol. 31, no. 11, pp. 1823-1833, April 1992.

[10]   David Casasent and Satoshi Ashizawa, "Synthetic aperture radar detection, recognition, and clutter rejection with new minimum noise and correlation energy filters," Optical Engineering, vol. 36, no. 10, pp. 2729-2736, October 1997.

[11]   Leslie M Novak, Gregory J. Owirka, and Christine M. Netishen, "Radar Target Identification using Spatial Matched Filters," Pattern Recognition, Vol. 27, no. 4, pp. 607-617, 1994.

[12]   Emanuel Marom and Hanni Inbar, "New Interpretations of Wiener Filters for Image Recognition," Journal of the Optical Society of America, Vol. 13, no. 7, pp. 1325-1330, July 1996.

[13]   Sovira Tan, Rupert C. D. Young, David M. Budgett, John D. Richardson, and Chris R. Chatwin, "A Pattern Recognition Wiener Filter for Realistic Clutter Backgrounds," Optics Communications, Vol. 172, no. 1, pp. 193-202, December 1999.

[14]   Athanasios Papoulis, *Probability, Random Variables, and Stochastic Processes*, Third Edition, McGraw-Hill, Inc, New York, 1991.

[15]   Merrill I. Skolnik, *Radar Handbook*, McGraw-Hill, Inc, New York, 1970.

[16]   E. Douglas Lynch, Russell D. Brown, David, D. Mokry, James M. VanDamme, and Michael C. Wicks, "Experimental HF Radar for Sub-surface Sensing," Seventh International Conference on Ground-Penetrating Radar, Lawrence, Kansas, May 1998.

[17]   Eugene Hecht, *Optics*, Third Edition, Addison-Wesley, Inc, New York, 1998.

[18]   John Parker Burg, *Maximum Entropy Spectral Analysis*, Ph.D. Thesis, Stanford University, May 1975.

[19]   J. Patrick Fitch, *Synthetic Aperture Radar*, Springer-Verlag, Inc, New York, 1988.

[20]   Rafael C. Gonzalez and Richard E. Woods, *Digital Image Processing*, Third Edition, Addison-Wesley Publishing Company, New York, 1992.

[21]   Sergios Theodoridis and Konstantinos Koutroumbas, *Pattern Recognition*, Academic Press, San Diego, California, 1999.

[22]   John G. Proakis and Dimitris G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, Third Edition, Prentice-Hall, Inc, Upper Saddle River, New Jersey, 1996.

[23]   Edwin E. Larson and Peter W. Birkeland, *Putnam's Geology*, Fourth Edition, Oxford University Press, New York, 1982.

[24]   James W. Jeter, III, *Three-dimensional Ray Tracing Simulation of a Synthetic Aperture Ground Penetrating Radar System*, Master Thesis, Rochester Institute of Technology, August 2002.

[25]   John Parker Burg, "Maximum Entropy Spectral Analysis," Thirty-Seventh Annual International SEG Meeting, Oklahoma City, Oklahoma, 1967.

[26]   L. C. Graham, "Synthetic Aperture Radar, Part 1 - System Principles," Goodyear Aerospace Corporation, 15 April 1970.

[27]   Khalid Sayood, *Introduction to Data Compression*, Second Edition, Academic Press, San Diego, California, 2000.

[28]   Fred E. Nathanson, J. Patrick Reilly, and Marvin N. Cohen, *Radar Design Principles*, Second Edition, McGraw-Hill, Inc, New York, 1991.

[29]   Rafael C. Gonzalez and Paul Wintz, *Digital Image Processing*, Second Edition, Addison-Wesley Publishing Company, New York, 1987.

[30]   J. V. DiFranco and W. L. Rubin, *Radar Detection*, Prentice-Hall, Inc, Englewood Cliffs, New Jersey, 1968.

[31]   Robert O. Harger, *Synthetic Aperture Radar Systems: Theory and Design*, Academic Press, New York, 1970.

[32]   Giorgio Franceschetti and Riccardo Lanari, *Synthetic Aperture Radar Processing*, CRC Press, LLC, Boca Raton, Florida, 1999.

[33]   Paul M. DeRusso, Rob J. Roy, Charles M. Close, and Alan A. Desrochers, *State Variables for Engineers*, Second Edition, John Wiley and Sons, Inc, New York, 1998.

[34]   Leon W. Couch II, *Digital and Analog Communication Systems*, Fourth Edition, Macmillan Publishing Company, New York, 1993.

[35]   Leland B. Jackson, *Digital Filters and Signal Processing*, Kluwer Academic Publishers, Boston, Massachusetts, 1986.

[36]   W. D. T, Davies, *System Identification for Self-Adaptive Control*, John Wiley and Sons, Inc, New York, 1970.

[37]   Russell D. Brown, Braham Himed, E. Douglas Lynch, and Michael C. Wicks, "Spatially Adaptive Ground Penetrating Radar," Eighth International Conference on Ground Penetrating Radar, Gold Coast, Australia, May 2000.

[38]   Russell D. Brown, Braham Himed, E. Douglas Lynch, and Michael C. Wicks, "Thinned Arrays for Ground Penetrating Radar," Eighth International Conference on Ground Penetrating Radar, Gold Coast, Australia, May 2000.

[39]   Ph. Refregier, "Filter Design for Optical Pattern Recognition: Multicriteria Optimization Approach," Optics Letters, Vol. 15, no. 15, pp. 854-856, August 1990.

[40]   Duane Hanselman and Bruce Littlefield, *Mastering MatLab 6: A Comprehensive Tutorial and Reference*, Prentice-Hall, Inc, Upper Saddle River, New Jersey, 2001.

[41]   Samir S. Soliman and Mandyam D. Srinath, *Continuous and Discrete Signals and Systems*, Second Edition, Prentice-Hall, Inc, Upper Saddle River, New Jersey, 1998.

[42]   Richard O. Duda and Peter E. Hart, *Pattern Classification and Scene Analysis*, John Wiley and Sons, Inc, New York, 1973.

[43]   Gary R. Olhoeft, "Ground Penetrating Radar: Use and Misuse," SPIE Conference on Subsurface Sensors and Applications, July 1999.

[44]   M. Lops, "Hybrid Clutter-map/L-CFAR Procedure for Clutter Rejection in Non-homogeneous Environment," IEE Proceedings – Radar, Sonar, and Navigation, Vol. 143, no. 4, pp. 239-245, August 1996.

[45]   Nicholas John Loy, *An Engineer's Guide to FIR Digital Filters*, Prentice-Hall, Inc, Upper Saddle River, New Jersey, 1988.

[46]   William M. Brown and Leonard J. Porcello, "An Introduction to Synthetic-Aperture Radar," IEEE Spectrum, Vol. 6, no. 9, pp. 52-62, September 1969.

[47]   David Knox Barton, *Handbook of Radar Measurement*, Prentice-Hall, Inc, Englewood Cliffs, New Jersey, 1969.

[48]   William A. C. Schmidt, "Modified Matched Filter for Cloud Clutter Suppression," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 12, no. 6, pp. 594-600, June 1990.

[49]   Jingsheng Sun and Roger A. Young, "Recognizing Surface Scattering in Ground-Penetrating Radar Data," Geophysics, Vol. 60, no. 5, pp. 1378-1385, September-October 1995.

[50]   Stacy Tantum and Leslie Collins, "ATR Algorithm Performance for the BRTRC Wichmann Ground Penetrating Radar System," UXO Forum 1999, Atlanta, Georgia, May 1999.

[51]   Remke L. Van Dam and Wolfgang Schlager, "Identifying Causes of Ground-Penetrating Radar Reflections Using Time-Domain Reflectometry and Sedimentological Analyses," Sendimentology, Vol. 47, no. 2, pp. 435-449, 2000.

[52]   Russell A. Kirsch, "Computer Determination of the Constituent Structure of Biological Images," Computers and Biomedical Research, Vol. 4, pp. 315-328, 1971.

[53]   Satoshi Takahashi and Susumu Miwa, "Suppression of Radar Clutter with a Weighted Moving Average Filter," Electronics and Communications in Japan, Vol. 77, no. 1, pp. 57-65, 1994.

[54]   P. R. Wu, "A Criterion for Radar Resolution Enhancement with Burg Algorithm," IEEE Transactions on Aerospace and Electronic Systems, Vol. 31, no. 3, pp. 897-915, 1995.

[55]   Chester R. Longwell and Richard F. Flint, *Introduction to Physical Geology*, Second Edition, John Wiley and Sons, Inc., New York, 1962.

[56]   Frank Press and Raymond Siever, *Earth*, W. H. Freeman and Company, San Francisco, 1974.

[57]   James Gilluly, Aaron C. Waters, and A. O. Woodford, *Principles of Geology*, Third Edition, W. H. Freeman and Company, San Francisco, 1968.